Matrix Calculations for Digital Photography

Paul M. Hubel and Jack Holm
HP Laboratories
Hewlett-Packard Company
Palo Alto, California 94304
USA

Graham D. Finlayson
Colour Research Group
University of Derby
Derby DE22 3BL
United Kingdom

Mark S. Drew
School of Computing Science
Simon Fraser University
Vancouver V5A 1S6
Canada

Introduction

Different digital cameras have different spectral sensitivities and opto-electronic conversion functions (OECF's), and therefore produce different data about the same scene. Correct interpretation of the data requires that it be presented in some sort of standard form. The most rigorous way to do this is to provide the raw data along with the camera color channel OECF's and spectral sensitivities. Another option is to transform the data itself into some sort of standard form with assumed OECF and spectral characteristics, or to specify a transformation. There are a number of paths to take in determining such transformations. The results obtained are variable, and depend on the scene and capture device spectral characteristics, and how the transformation to the standard data form was determined.

Since the spectral spaces spanned by different cameras are different, and the intended use of the data in digital photography is to produce a reproduction for viewing, the obvious choice for a standard spectral space is a spectral space spanned by a set of color matching functions, *e.g.* a color space. This paper outlines several methods for determining 3x3 matrix transformations from camera spectral spaces to a standard color space based on the ITU-R BT.709 red, green, and blue (RGB) primaries.[1] Since these transformations are intended for digital cameras, they are based on either of two assumptions about the scene spectral correlation statistics: standard surface reflection statistics as represented by a Macbeth Color Checker (MBCC),[2,3] or the assumption of no spectral correlation (maximum ignorance).[4]

This paper only deals with 3x3 matrix transformations because such transformations are most appropriate when the relationship between the scene radiance and the radiance incident on the sensor is variable, different illumination sources are used, and the colorants found in the scene are unknown or highly variable. If the illumination source, colorants, and corresponding spectral correlation statistics of the scene or original are known, as is typically the case with film scans, the actual correlation statistics should be applied by using the appropriate spectral correlation matrix, or through the use of a characterization target with similar spectral correlation statistics. In such cases greater transformation accuracy may be obtained using nonlinear transformations, such as polynomials or multi-dimensional look-up-tables. Guidance as to which procedures to use in determining transformations for digital cameras was taken from a preliminary proposal for a new ISO standard in this area,[5] with emphasis placed on white point preserving transformations.

The transformation determination methods outlined in this paper were used to obtain transformations for an actual digital camera. These transformations were then applied to image data obtained from the camera. Images of test charts and pictorial scenes were processed, and CIE L*a*b* ΔE [6] and CMC ΔE [7] values determined for the test chart patches. A preferred reproduction model[8] was also applied to the transformed data, and the results of this additional processing step evaluated similarly. The preferred rendition data was then printed using a calibrated printer, and a preliminary subjective evaluation conducted.

These experiments indicate that the white point preserving transformation based on the maximum ignorance assumption produces visual results which are as good or better than those produced by transformations based on surface reflectance spectral correlation statistics. This result is particularly significant in view of the large amount of variability which occurs when trying to determine transformations from images of test charts, the ease with which the white point preserving maximum ignorance (WPPMI) transformation can be determined, and the fact that only one 3x3 transformation matrix is used, as opposed to a different one for every type of scene illumination. Of particular interest is the fact that the WPPMI transformation produced results which were significantly better than the other transformations when the capture illumination chromaticity was very different from the rendering white point chromaticity. In effect, a chromatic adaptation transform based on semi-arbitrary sensor RGB spectral sensitivities outperformed one based on spectrally sharpened cone sensitivities (for an imaging application). Another interesting result was the rather limited correlation between the ΔE metrics and color reproduction quality.

Discussion

Digital cameras and scanners typically do not capture information about scenes or original hardcopy that is in the most strict sense colorimetric. The spectral responses and distribution of the ρ, β, and γ cones of the eye are most likely the result of evolutionary factors, and are close to optimal considering the chromatic aberration of the eye lens. Even if one feels that products such as digital cameras also evolve, the fact that virtually all such devices possess achromatic optics results in different driving forces in sensor design. The use of achromatic optics for digital capture is dictated by the fact that image data is used to produce reproductions for viewing. These reproductions are viewed at different magnifications and distances, making it impossible to know what the human visual system (HVS) spatial sampling will be for each color of the reproduction when the image is captured.

Signal-to-noise considerations also preclude the use of linear combinations of the HVS cone responses. Most current sensors are hard pressed to equal the extremely large dynamic range capability of the HVS even without channel differencing. As sensor technology improves, the dynamic range requirements may be more easily met, but cost factors will continue to place practical limits on dynamic range. At the limit, even if silicon becomes virtually perfect and free, photon noise statistics will continue to make channel differencing undesirable from a signal-to-noise standpoint. The large amount of overlap in the spectral sensitivities of the ρ and γ cones requires large amounts of differencing (or ratioing) with any all-positive linear combination of these sensitivities. It is more practical to try to capture image data that is representative of the HVS color signals after differencing.

Finally, even if colorimetric capture were desirable, it is practically difficult to achieve. Film responses have long been limited by sensitizing dye characteristics. Different materials are used to modulate the spectral response of the color channels in digital cameras, but real and practical materials must still be used. It is also important to remember that actual cone sensitivities are not known with absolute precision, and vary between individuals to some extent. CIE XYZ based color matching functions are a formally standardized best guess at mean sensitivities.[9]

The above discussion argues against the use of cone responses for image capture, however, this argument *does not* extend to the description of the scene or original captured. It *is* necessary to be able to determine, as accurately as possible, the stimulus the scene or original would have created in the HVS. The purpose of the preceding discussion is to provide support for the fact that the results presented in this paper all relate to the transformation of the data captured by the sensor into colorimetrically defined RGB data.

An important side note is that it is sometimes necessary to go through intermediate stages in producing a color reproduction. The form of the data in these intermediate stages may not have meaning in terms of the HVS. For example, a film scanner may have RGB analysis, but the spectral response of such a scanner should be optimized to produce the most accurate measurements of the film dyes. These measurements are then transformed into colorimetric information about the intended reproduction based on the film dye spectral characteristics. Since color film transparencies are intended for direct viewing, their dyes modulate the RGB light in visually meaningful ways, so analysis requires scanner RGB sensitivities that also tend to be visually meaningful. On the other hand, color film negatives are meant to modulate the RGB light as seen by photographic paper. The red sensitivity of photographic paper peaks at a much longer wavelength than that of the HVS. It would therefore be inappropriate to measure color negatives using an HVS based RGB analysis; this is in addition to the fact that the image information is in negative form.

Experimental Outline

Two basic transformation methods were explored in the experiments presented here: least squares (LS) regressions in a standard linear RGB color space of image data captured using a MBCC under different illumination conditions, and WPPMI transformations based on camera spectral responses. The advantage of the white point preservation constraint when assuming maximum ignorance (MI) spectral correlation statistics is discussed elsewhere.[10] For consistency's sake, the least squares regressions were performed on linear RGB values. There may be some advantage to minimizing the mean square error in a more perceptual RGB space, such as one with a gamma function of approximately 2.2, but this was not investigated. However, it is important to note that the mean square error was minimized in linear ITU-R BT.709 RGB space, as opposed to linear CIE XYZ space.

Three methods were employed to determine LS transformations: simple LS regressions, white point preserving LS regressions (WPPLS), and weighted white point preserving LS regressions (WTWPPLS). With the WTWPPLS regressions, the weights of selected patches of the MBCC were increased to produce the most visually pleasing result when applying the transformation to image data. These weights may depend to some extent on the spectral response characteristics of the camera for which the transformation is determined.

Camera Data Linearization

The first step in determining the desired transformations is to linearize the data obtained from the camera with respect to the source or scene radiance. This linearization is more difficult with the chart based LS methods, because some of the non-linearity is due to camera flare light which is scene dependent. The spectral response measurements required to determine the WPPMI transformation must also be accurate, but can generally be focal plane measurements which do not change.

In these experiments, the camera data was linearized by measuring the camera and focal plane OECF's[11] of the camera for a variety of scene dynamic ranges and mean reflectances, and constructing a flare model which predicts the camera flare based on focal plane image statistics. With the images of the MBCC, it is also possible to estimate the amount of camera flare light by manually determining a value that produces linear data for the six gray scale patches when the estimated amount is removed. Accurate OECF measurements for camera data linearization are extremely important, particularly with the LS methods, since the regression tries to transform the chart image data to aim linear values.

Scene Illumination Conditions

The following illumination conditions were studied by capturing images of MBCC's and real scenes with people and other natural objects:

1. Natural daylight outdoors (direct sunlight plus skylight) mid-afternoon on a clear day; measured correlated color temperature (CCT) 5352K.
2. Typical quartz-halogen lighting; measured CCT 3078K.
3. Fluorescent office lighting; measured CCT 3254K.

The following additional illumination conditions were studied by capturing images of MBCC's only in a Macbeth SpectraLight viewing booth with the following illumination settings selected:

4. CIE Illuminant D_{65} simulator (filtered tungsten); measured CCT 6419K.
5. CIE Illuminant D_{50} simulator (filtered tungsten); measured CCT 4937K.
6. Viewing booth cool white fluorescent; measured CCT 4038K.
7. Viewing booth horizon illumination simulator; measured CCT 2298K.

Since three LS transformation methods were applied to seven illumination conditions, a total of 21 LS transformation matrices were obtained (see table 2).

Determination of Aim RGB Values for LS Regressions

The aim linear RGB values for each patch of the MBCC were determined by measuring the MBCC under each illumination condition using a Photo Research PR650 spectroradiometer. CIE XYZ values were then calculated from the spectral measurements. These values were transformed to XYZ values appropriate for the ITU-R BT.709 CIE illuminant D_{65} white point using the sharp transform described by Finlayson and Drew.[12] The D_{65} XYZ values were then transformed into the aim RGB values using the appropriate matrix for the ITU-R BT.709 RGB primaries.

Recent research indicates that spectrally sharpened transformations do a better job of accounting for different illumination adaptation conditions than von Kries transformations.[13] The CIECAM97 appearance model[14] makes use of Bradford transformations, another type of sharp transformation. One could argue that it would have been better to use the actual Bradford transformations for this work, but to do so would have required additional software development. The differences between the sharp transformations used and the Bradford transformations are not too large, so the results produced by each should be similar, so long as the same transformation method is used in determining the aim values and in processing the camera data. In any case, it seems likely that illumination adaptation transformations will continue to evolve, at least until a physiological explanation of the improved performance of the spectrally sharpened transforms is found.

Determination of the WPPMI Transformation

In the case of the WPPMI transformation, only one 3x3 matrix is used for all conditions. Since no spectral correlation is assumed, it is not necessary to consider the effect of the illumination source on the spectral correlation statistics. The method for determining WPPMI transformations is outlined in the appendix of the paper "White-Point Preserving Color Correction" by two of the authors of this paper, which can also be found in these proceedings.[10] The actual WPPMI transformation matrix used in this study is provided in table 1.

Preferred Reproduction Processing

It is reasonably well known that preferred reproductions of pictorial scenes do not duplicate the scene colorimetrically. Even if it is possible to obtain an exact colorimetric description of the scene, the most pleasing reproduction, if measured, will usually have different colorimetric values. Some, but not all of the difference can be accounted for by the fact that the reproduction may be viewed under different conditions than the original scene. Other differences are attributable to the characteristics of the reproduction medium, and viewer preferences.

In order to simulate a real imaging chain, it is therefore necessary to apply some sort of preferred rendering algorithm to the image data. If the scene RGB estimates obtained using a particular transformation are rendered unacceptably by the preferred rendering algorithm, the two are not compatible, and either a different transform, or a different rendering algorithm must be used. In this study, the preferred reproduction algorithm developed previously by one of the authors is used.[8]

Preliminary Subjective Evaluation

Ultimately, the results of transformations used for digital photography will be evaluated by the end user. In the experiments described here, a small subjective study was conducted to provide preliminary information in this area. The prints for the subjective study were produced by applying the different transformations to the image data of the MBCC. The transformed standard RGB data was then processed using the preferred reproduction model, and printed on a calibrated printer. These prints were divided into sets according to the illumination source under which the image was captured; four prints for each source resulting from the four transformation determination methods. The prints were labeled in a location not visible to the observers.

Four observers were then asked to rank the prints in order of preference. The observers were provided with the MBCC illuminated using the original source for the captured images. The reproductions were viewed using the same source in the case of the studio tungsten and office fluorescent images, in natural daylight in the case of the D_{50} images, and in office fluorescent in the case of the D_{65}, Booth CWF, and Booth Horizon images. Successive evaluations with adaptation time were encouraged where the reproduction and original viewing sources were different, although this was not enforced. The conditions of the study were meant to be similar to those used in evaluating photographic reproductions in general.

The ranks assigned by the observers to each print were then averaged to produce a visual ranking score. Note that lower scores indicate preferred reproductions, and scores which are closer to the limits of 1 and 4 indicate that the preferences expressed were consistent among the observers.

Results

The numerical results of the experiments conducted are presented in tables 1 through 4.

Table 1: WPPMI Transformation Matrix

1.3203	-0.2063	-0.1162
-0.0842	1.3240	-0.2388
0.0188	-0.3927	1.3714

Table 2: LS Transformation Matrices

Transformation Method/ Illumination	LS	WPPLS	WTWPPLS
Daylight	2.825 0.185 -0.347 -0.208 1.145 -0.126 0.124 -0.386 1.588	1.026 0.093 -0.114 -0.081 1.155 -0.071 0.066 -0.267 1.193	1.049 0.129 -0.170 -0.088 1.182 -0.090 0.030 -0.306 1.264
Tungsten	1.475 0.230 -0.600 -0.187 1.052 -0.099 0.095 -0.464 2.745	1.027 0.100 -0.121 -0.124 1.177 -0.052 0.070 -0.293 1.214	1.053 0.140 -0.184 -0.130 1.199 -0.068 0.034 -0.355 1.308
Fluorescent	1.897 0.040 -0.377 -0.161 1.000 -0.094 0.070 -0.310 2.494	1.123 -0.039 -0.083 -0.094 1.147 -0.053 0.063 -0.246 1.182	1.138 0.005 -0.141 -0.090 1.156 -0.065 0.011 -0.229 1.216
Booth D65	3.028 0.117 -0.270 -0.265 1.095 -0.164 0.051 -0.328 1.276	1.141 -0.045 -0.093 -0.083 1.173 -0.088 0.010 -0.164 1.149	1.177 0.003 -0.174 -0.065 1.214 -0.145 -0.017 -0.213 1.223
Booth D50	2.489 0.162 -0.364 -0.223 1.081 -0.154 0.064 -0.365 1.602	1.135 -0.018 -0.114 -0.087 1.174 -0.084 0.026 -0.201 1.170	1.182 0.029 -0.205 -0.065 1.198 -0.130 -0.002 -0.256 1.250
Booth CWF	2.542 0.241 -0.357 -0.213 0.984 -0.156 0.118 -0.323 2.036	0.953 0.154 -0.102 -0.058 1.125 -0.064 0.016 -0.186 1.163	0.961 0.208 -0.162 -0.045 1.156 -0.106 -0.003 -0.215 1.209
Booth Horizon	1.117 0.263 -1.056 -0.183 1.201 -0.071 0.196 -0.753 4.935	1.083 0.168 -0.245 -0.158 1.249 -0.092 0.199 -0.664 1.458	1.134 0.178 -0.303 -0.143 1.262 -0.119 0.178 -0.692 1.506

Table 3: CIE L*a*b* and CMC ΔE's

Transformation Method/ Illumination	LS	WPPLS	WTWPPLS	WPPMI
Natural Daylight	CIE L*a*b* ΔE Mean = 3.7, Max = 7.9 CMC ΔE Mean = 2.8, Max = 5.7	CIE L*a*b* ΔE Mean = 6.6, Max = 17.2 CMC ΔE Mean = 4.0, Max = 6.6	CIE L*a*b* ΔE Mean = 5.1, Max = 10.6 CMC ΔE Mean = 3.4, Max = 5.3	CIE L*a*b* ΔE Mean = 6.4, Max = 15.6 CMC ΔE Mean = 4.3, Max = 8.0
Studio Tungsten	CIE L*a*b* ΔE Mean = 3.6, Max = 7.2 CMC ΔE Mean = 2.8, Max = 5.2	CIE L*a*b* ΔE Mean = 6.7, Max = 18.4 CMC ΔE Mean = 4.1, Max = 6.8	CIE L*a*b* ΔE Mean = 5.4, Max = 9.2 CMC ΔE Mean = 3.7, Max = 6.5	CIE L*a*b* ΔE Mean = 7.0, Max = 16.5 CMC ΔE Mean = 5.4, Max = 11.2
Office Fluorescent	CIE L*a*b* ΔE Mean = 3.3, Max = 7.9 CMC ΔE Mean = 2.4, Max = 4.6	CIE L*a*b* ΔE Mean = 5.4, Max = 16.2 CMC ΔE Mean = 3.3, Max = 6.3	CIE L*a*b* ΔE Mean = 4.5, Max = 10.7 CMC ΔE Mean = 3.0, Max = 5.7	CIE L*a*b* ΔE Mean = 6.4, Max = 19.3 CMC ΔE Mean = 4.4, Max = 12.3
Macbeth SpectraLight Booth D65	CIE L*a*b* ΔE Mean = 3.6, Max = 8.2 CMC ΔE Mean = 2.7, Max = 4.7	CIE L*a*b* ΔE Mean = 7.2, Max = 18.6 CMC ΔE Mean = 4.2, Max = 7.4	CIE L*a*b* ΔE Mean = 5.9, Max = 11.7 CMC ΔE Mean = 3.9, Max = 6.0	CIE L*a*b* ΔE Mean = 4.6, Max = 10.6 CMC ΔE Mean = 3.4, Max = 6.3
Macbeth SpectraLight Booth D50	CIE L*a*b* ΔE Mean = 3.4, Max = 7.4 CMC ΔE Mean = 2.6, Max = 5.0	CIE L*a*b* ΔE Mean = 6.9, Max = 18.6 CMC ΔE Mean = 3.9, Max = 6.5	CIE L*a*b* ΔE Mean = 5.4, Max = 9.9 CMC ΔE Mean = 3.5, Max = 5.5	CIE L*a*b* ΔE Mean = 4.7, Max = 11.2 CMC ΔE Mean = 3.6, Max = 6.2
Macbeth SpectraLight Booth CWF	CIE L*a*b* ΔE Mean = 3.2, Max = 5.9 CMC ΔE Mean = 2.5, Max = 6.8	CIE L*a*b* ΔE Mean = 6.8, Max = 13.9 CMC ΔE Mean = 3.9, Max = 5.8	CIE L*a*b* ΔE Mean = 5.5, Max = 10.7 CMC ΔE Mean = 3.6, Max = 5.9	CIE L*a*b* ΔE Mean = 6.1, Max = 16.0 CMC ΔE Mean = 4.3, Max = 9.4
Macbeth SpectraLight Booth Horizon	CIE L*a*b* ΔE Mean = 4.3, Max = 10.5 CMC ΔE Mean = 3.0, Max = 6.0	CIE L*a*b* ΔE Mean = 4.9, Max = 12.0 CMC ΔE Mean = 3.4, Max = 6.5	CIE L*a*b* ΔE Mean = 4.6, Max = 9.7 CMC ΔE Mean = 3.5, Max = 6.8	CIE L*a*b* ΔE Mean = 9.4, Max = 21.8 CMC ΔE Mean = 6.6, Max = 13.1

Table 4: Preferred Reproduction ΔE's and Visual Rankings (lower numbers are better for rankings).

Transformation Method/ Illumination	LS	WPPLS	WTWPPLS	WPPMI
Studio Tungsten	CIE L*a*b* ΔE Mean = 10.9, Max = 19.0 CMC ΔE Mean = 7.2, Max = 21.9 Visual Ranking = 3.5	CIE L*a*b* ΔE Mean = 8.5, Max = 18.5 CMC ΔE Mean = 6.8, Max = 21.9 Visual Ranking = 1.75	CIE L*a*b* ΔE Mean = 10.2, Max = 24.6 CMC ΔE Mean = 7.5, Max = 21.9 Visual Ranking = 2.75	CIE L*a*b* ΔE Mean = 15.4, Max = 34.3 CMC ΔE Mean = 10.2, Max = 21.6 Visual Ranking = 2.0
Office Fluorescent	CIE L*a*b* ΔE Mean = 11.1, Max = 19.4 CMC ΔE Mean = 7.2, Max = 19.0 Visual Ranking = 3.0	CIE L*a*b* ΔE Mean = 9.4, Max = 18.5 CMC ΔE Mean = 6.8, Max = 19.5 Visual Ranking = 2.75	CIE L*a*b* ΔE Mean = 10.4, Max = 23.5 CMC ΔE Mean = 7.1, Max = 19.5 Visual Ranking = 2.75	CIE L*a*b* ΔE Mean = 16.5, Max = 35.3 CMC ΔE Mean = 10.1, Max = 19.5 Visual Ranking = 1.5
Macbeth SpectraLight Booth D65	CIE L*a*b* ΔE Mean = 12.5, Max = 24.4 CMC ΔE Mean = 7.7, Max = 15.2 Visual Ranking = 2.0	CIE L*a*b* ΔE Mean = 9.1, Max = 26.6 CMC ΔE Mean = 6.6, Max = 15.1 Visual Ranking = 2.75	CIE L*a*b* ΔE Mean = 11.0, Max = 30.6 CMC ΔE Mean = 7.4, Max = 15.8 Visual Ranking = 2.5	CIE L*a*b* ΔE Mean = 14.5, Max = 31.3 CMC ΔE Mean = 8.9, Max = 16.4 Visual Ranking = 2.75
Macbeth SpectraLight Booth D50	CIE L*a*b* ΔE Mean = 11.8, Max = 21.2 CMC ΔE Mean = 7.3, Max = 13.3 Visual Ranking = 3.0	CIE L*a*b* ΔE Mean = 8.8, Max = 22.7 CMC ΔE Mean = 6.4, Max = 13.9 Visual Ranking = 2.25	CIE L*a*b* ΔE Mean = 10.8, Max = 27.5 CMC ΔE Mean = 7.2, Max = 14.7 Visual Ranking = 3.0	CIE L*a*b* ΔE Mean = 14.6, Max = 29.6 CMC ΔE Mean = 8.7, Max = 16.0 Visual Ranking = 1.75
Macbeth SpectraLight Booth CWF	CIE L*a*b* ΔE Mean = 12.7, Max = 20.7 CMC ΔE Mean = 8.1, Max = 14.1 Visual Ranking = 2.75	CIE L*a*b* ΔE Mean = 9.7, Max = 23.8 CMC ΔE Mean = 7.3, Max = 14.1 Visual Ranking = 2.75	CIE L*a*b* ΔE Mean = 11.2, Max = 25.7 CMC ΔE Mean = 8.0, Max = 14.1 Visual Ranking = 2.5	CIE L*a*b* ΔE Mean = 17.6, Max = 36.4 CMC ΔE Mean = 10.8, Max = 20.0 Visual Ranking = 2.0
Macbeth SpectraLight Booth Horizon	CIE L*a*b* ΔE Mean = 11.6, Max = 18.0 CMC ΔE Mean = 7.4, Max = 14.6 Visual Ranking = 2.5	CIE L*a*b* ΔE Mean = 11.8, Max = 17.9 CMC ΔE Mean = 7.9, Max = 15.3 Visual Ranking = 2.25	CIE L*a*b* ΔE Mean = 13.1, Max = 21.2 CMC ΔE Mean = 8.5, Max = 15.3 Visual Ranking = 3.5	CIE L*a*b* ΔE Mean = 17.0, Max = 41.0 CMC ΔE Mean = 11.3, Max = 22.2 Visual Ranking = 1.75

Note: The presence of objects other than the MBCC in the natural daylight scene resulted in the preferred reproduction algorithm producing different results for this image. It is not included in the above comparison because of this difference in processing.

Possibly the most interesting result is that the WPPMI based reproductions had the highest ΔE's, but were also consistently preferred by the observers in all cases except the D_{65} case, where there is no change in white point between capture and viewing.

Conclusions

The experience gained in conducting the described experiments, in combination with the numerical results presented, suggests the following conclusions:

1. Any non-visible radiation to which the camera is sensitive must be completely blocked so that it produces no measurable change in the data output by the camera.
2. The actual camera OECF must be well known for the calculation of reliable transformations. This includes the focal plane OECF for each channel and the image specific contribution of flare light in the camera.
3. The measurements obtained from the MBCC, both using a camera and a telescopic spectroradiometer, were dependent on the illumination geometry. Transformations determined based on these measurements are therefore variable. However, the MBCC appeared to have a significant specular reflection component under strong directional illumination.

Charts with more Lambertian surfaces may produce better results. A transmissive test target could also improve the reliability of target based transformations, but might introduce different spectral correlation statistics.

4. The spectral sensitivity measurements on which the WPPMI transformation is based seemed to be easier to repeat with greater accuracy than the test target measurements. This is probably because the test target measurements were so dependent on illumination geometry.

5. Relative errors between colors in a scene tend to be more noticeable than absolute errors, particularly when they are in different directions. The WPPMI and WTWPPLS transformations tended to produce the fewest objectionable relative errors. In the case of the WTWPPLS transformation, this is probably because the weights determined are based on visual evaluation.

6. The WPPMI transformation tends to produce somewhat more color saturation than any of the LS transformations. This is probably because the surface color spectral correlation statistics approximated by the MBCC are for the surface color gamut, which does not include extremely saturated (*i.e.* spectral) colors. However, the saturation produced by the WPPMI transformation is not equivalent to a LS transformation concatenated with a saturation boost.

7. Since preferred reproduction algorithms tend to increase saturation, it is particularly important to reproduce the neutral scale correctly and minimize relative color errors in transformations which are to be followed by preferred reproduction processing.

8. The single WPPMI matrix in combination with channel balancing multipliers seemed to do a better job of dealing with extreme illumination sources, such as the horizon illumination.

9. It is difficult to gauge the acceptability of color reproduction by viewing reproductions of MBCC patches. Errors which do not appear large when viewing the patches may be unacceptable in reproductions of real scenes. Conversely, some types of errors in patch reproduction are quite acceptable in reproductions of real scenes.

10. Neither the CIE L*a*b* ΔE nor the CMC ΔE metrics are reliable predictors of the visual objectionability of color errors in digital photography, particularly if only mean and maximum values are used.

11. Overall, the WPPMI transformation is the easiest to reliably determine and produces visual results which are as good or better than those obtained using transformations determined by other methods, especially when preferred reproduction processing is applied.

References

1. ITU-R BT.709 - 1993, *Basic parameter values for the HDTV standard for the studio and for international programme exchange*.
2. C.S. McCamy, H. Marcus, & J.G. Davidson, "A Color Rendition Chart," *J. Appl. Photog. Eng.*, **2**:95-99, 1976.
3. G. D. Finlayson & M.S. Drew, “Constrained Least-Squares Regression in Color Space,” to be published in the *Journal of Electronic Imaging*, 1997.
4. G.D. Finlayson & M.S. Drew, "The Maximum Ignorance Assumption with Positivity," *Proceedings, IS&T/SID Fourth Color Imaging Conference: Color Science, Systems, and Applications*, p. 202-205, 1996.
5. ISO TC42 (Photography) and TC130 (Graphic Technology) Joint New Work Item Proposal: *Digital still picture cameras - Methods for the transformation of sensor data into standard colour spaces*, February 1997.
6. CIE Publication No. 15.2 - 1986, *Colorimetry, 2nd edition.*
7. M.H. Brill, “Suggested Modification of CMC Formula for Acceptability,” *Color Res. and App.* **17**(6):402-404, 1994.
8. J. Holm, “Photographic Tone and Colour Reproduction Goals,” *Proceedings, CIE Expert Symposium ‘96: Colour Standards for Image Technology*, p. 51-56, 1996.
9. G. Wyszecki & W.S. Stiles, *Color Science: Concepts and Methods, Quantitative Data and Formulae*, Wiley, 1982.
10. G. D. Finlayson & M.S. Drew, “White-Point Preserving Color Correction,” these proceedings.
11. ISO/DIS 14524, *Photography - Electronic still picture cameras - Methods for measuring opto-electronic conversion functions (OECF's)*, 28 May 1997.
12. G.D. Finlayson, M.S. Drew, & B.V. Funt, “Spectral Sharpening: Sensor Transformations for Improved Color Constancy,” *JOSA*, **5**:1,553-1,563 (1994).
13. P.M. Hubel & G.D. Finlayson, “Sharp Transformations for Color Appearance,” to be published in the proceedings of the SPIE Color Imaging: Device Independent Color, Color Hardcopy, and Graphic Arts III Conference, 1998.
14. R.W.G. Hunt & M.R. Luo, “The Structure of the CIE 1997 Colour Appearance Model (CIECAM97),” manuscript handed out at the AIC Conference, 1997.