COLOR CORRECTION AND INTERPOLATION WITH FEW SAMPLES

H. J. Trussell and Huiwen Zeng

Dept. of Electrical and Computer Engineering
North Carolina State University
Raleigh, NC 27695-7911

M. J. Vrhel

ViewAhead Technology
18200 NE Union Hill Road
Redmond, WA 98052

ABSTRACT

With the proliferation of digital cameras, more consumers are faced with the problem of making color corrections to their pictures. while most picture editors allow some white point, or illuminant, correction and some enhancements, these methods work on the image as a whole rather than on specific regions or colors. This work describes a simple method, that allows the user the ability to correct specific colors in specific locations in a picture. It is shown that relatively few sample colors are needed to produce a well-behaved interpolation in three-dimensional color spaces. A quantitative measure of performance is obtained by using the method to correct for illumination.

1. INTRODUCTION

Since digital cameras are becoming commonplace, more people are uploading pictures from their cameras, reviewing them on personal computers and editing them for future display. A common problem that occurs is that the colors of the recorded image do not match those of the object that was photographed. There is much work done on estimating the illumination of the scene and, in turn, estimating the tristimulus values, e.g., RGB, XYZ, CIELab, of the objects. However, the user can often see immediately upon looking at the results that these estimates are inaccurate. The problem is usually caused by failure of the camera or user to determine accurately the scene conditions. In any case, the user knows the approximate color of an object in the picture and recognizes that the recorded color is not correct. The problem that this paper addresses is to make a correction that is accurate and requires minimal effort by the user.

User-based color adjustments in most commercial software packages are performed by transforming the image into a color space in which the three color components consist of luminance, hue, and saturation [1]. The user is then allowed to globally adjust the values of the image along these components. With this approach, it is difficult to use knowledge of known object colors in the image and to optimize the controls to achieve these colors.

In the following approach, given a set of known colors that need to be adjusted, it is easy to generate an interpolating function. Evaluating the performance of the function is much more difficult. The obvious way to measure performance is by user observation. However, this is subjective and, although appropriate for the majority of uses of this method, is not likely to convince critics. We propose to test the method on an illuminant correction problem where the target image is known. If the method is shown to perform well in this quantitative case, it is likely to perform reasonably well in the subjective cases.

2. THE PROBLEM

The user has obtained image data in the form of tristimulus values associated with each pixel in the image. The tristimulus values may be recorded as sRGB, RGB, L^*, a^*, b^*, depending on the camera and the user-defined settings of that camera. The user recognizes that certain colors in the image are not accurate and wishes to correct them. The problem is that the user does not know the tristimulus values of the correct color. However, he or she can usually chose the correct color from an appropriate palette of colors.

It is easy to map one tristimulus vector onto another. The problem is complicated by the fact that if a single color is changed, then other colors must also be changed to produce a continuous transformation. Furthermore, it is impossible for the user to specify all of the colors that should be changed in a pictorial scene where shading, contours and shadows affect the specific values and the appearance of the colors of the objects.

The basic problem is one of creating a function from one three-dimensional color space to another. Given a set of tristimulus values in an original picture, $\{\mathbf{o}_k\}$ and a set of target values in the enhanced picture, $\{\mathbf{t}_k\}$, we wish to find a continuous mapping, $\mathcal{F}(\mathbf{o}_k) = \mathbf{t}_k$, such that the colors that are not specified in the original picture are mapped to reasonable values in the target picture.

The function could take on a number of forms. For example, the function could be defined by using a low dimensional look-up table (LUT) combined with three dimen-

sional interpolation. This method is commonly used in the definition of mappings in printer and scanner characterization [2]. In the calibration scenario, where extreme accuracy is desired, the LUT is obtained by obtaining measurements from a grid of samples. In most cases, the grid is a minimum of $5 \times 5 \times 5$. While this implies only 125 samples to be specified, this is far too many to expect a user to input by hand. To be an effective consumer tool, we would hope to produce a reasonable function with fewer than 10 samples. As a simple starting point, we use low order polynomial mappings to test the feasibility of using a few points to determine a useful mapping. For example, four points can be used to create an interpolation mapping of the form,

$$\mathcal{F}(\mathbf{o}_k) = \mathbf{M}\mathbf{o}_k + \mathbf{b} = \mathbf{t}_k \qquad (1)$$

where $\mathbf{M}$ is a 3x3 matrix and $\mathbf{b}$ is a 3-vector. More than four linearly independent points can be used to create least squares mappings, or nonlinear terms can be added to allow exact matching.

Because we will need to control the definition of the function over a wide range of values, we will design our definition to allow fixed points, as well as changed points. It is suspected that there could be cases where the user defines the desired changes in a very small region of the color space. The creation of a global function based on such a range of definition could lead to unstable behavior. By allowing easy definition of pixels that retain their original values, the stability of the function can be assured.

3. THE EXPERIMENTAL SETUP

A digital image is a candidate for testing the method in its subjective enhancement mode. To test the method in a quantitative mode of illuminant correction, we need special image data. We will use the 31-band image described in [3]. Since this image is defined by its reflectance at each pixel, it is easy to generate the tristimulus values that are obtained under any illuminant for which the spectrum is defined. The equation for the tristimulus values under illuminant $\mathbf{L}_1$ is given by

$$\mathbf{t}_1 = \mathbf{A}^T \mathbf{L}_1 \mathbf{r}$$

where $\mathbf{r}$ is the 31-band reflectance vector, $\mathbf{L}_1$ is the diagonal matrix that defines the illuminant, and $\mathbf{A}$ is the 31×3 matrix defined by the CIE color matching functions. The tristimulus values can be transformed by any of the many color transformations to give data that can stored for later use. We will use the sRGB space to represent the image. This image is shown in Fig. 1 using the standard uniform illuminant. This image is defined as the target image.

For our experiment, a second image is generated using a bluish illuminant shown in Figure 2 . The image is shown in Fig. 3, and shows a distinct bluish cast. This image is the
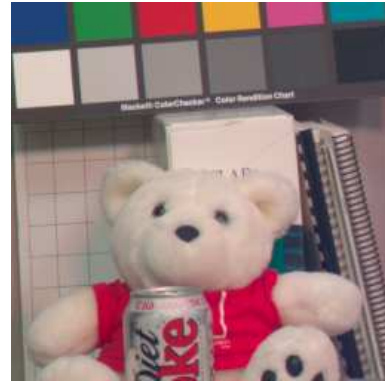

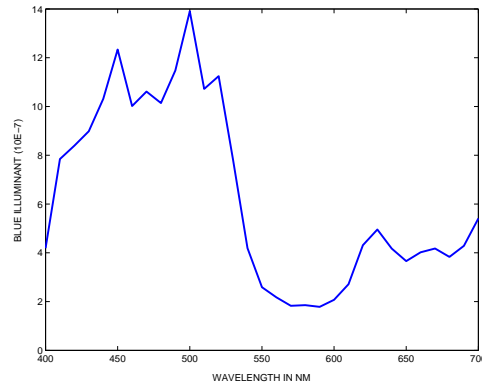**Fig. 1**. Bear image under uniform illumination

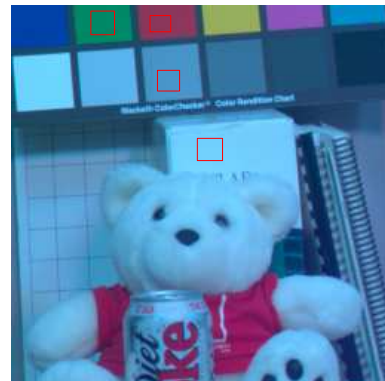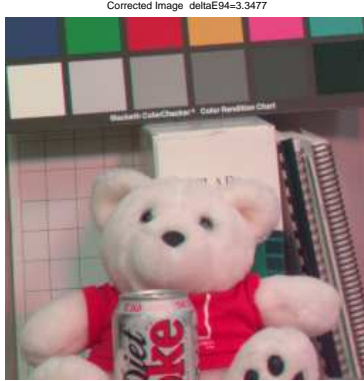**Fig. 2**. Spectral distribution of blue illuminant

Fig. 3. Bear image under illumination given in Figure 2. Red squares are regions selected for correction

Fig. 4. Bear image after four point correction

Fig. 5. Binary error image. White regions have errors larger than 4

original image that should be changed to look like the target image.

The problem is to define a number of pixels in the original image and create a mapping that takes the values of those pixels to the values of the corresponding points in the target image. The definitions of the regions to be moved are obtained from the user by placement of the cursor to create a rectangle in the original image (as shown by the red rectangles in Figure 3). We use the average value of the sRGB values in the original and target images to define the mapping. For colors that remain the same, the same method is used but only the original image is relevant.

To use the same method for color enhancement, the target image is replaced by a color palette. The color in the original image is defined by the average of the pixels in a rectangle that is drawn using the cursor. An additional step is needed to define the target color in the palette by using the cursor a second time. For this testing, we will also use a scanned image of a faded photograph, shown in Fig. 7.

4. RESULTS AND EXAMPLES

From experience with color mappings related to characterization of various devices, we were aware of the importance of the neutral axis in color spaces. Thus, we guessed that any definition should include points near the neutral axis. Some of the most obvious color errors appear in the more saturated regions of the color space. Furthermore, if colors near the boundary of the color gamut are not defined, there is a danger of generating an unstable mapping. Thus, we believed that using saturated colors was necessary in the definition of the mapping.

The surprising result is that very good illuminant correction can be achieved with as few as 4 samples. Using samples in the red and green regions and 2 samples in the neutral tones, produced visually very good results shown in Figure 4, with an average $\Delta E94 = 3.35$. In this example,
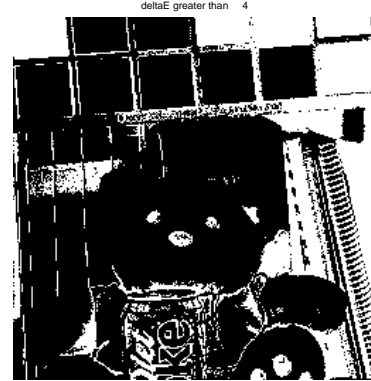
the affine mapping given in Equation 1 was used.

An examination of the error showed the largest color errors to be around edges, which would not be readily visible to an observer. A binary image showing those errors larger than 4.0 is shown in Fig. 5. Results that showed poor performance included those that failed to define samples either in the saturated regions or in the neutral region. In addition, fixing samples in the dark regions led to very high errors in these regions. This is a function of the definition of the CIELab color space.

As mentioned, additional samples can be used to improve the results and nonlinear terms can be introduced. Table 1 displays results of correcting the image shown in Figure 6 from the same blue illuminant. Figure 6 is a simulated image of 64 Munsell colors that have been desaturated to avoid sRGB gamut problems under the uniform and blue illuminations. The method column in Table 1 refers to the equation used to perform the adjustment. For example, Affine + RG implies the use of a transformation of the form

$$\mathcal{F}(\mathbf{o}_k) = \mathbf{t}_k = \mathbf{M}\mathbf{z}_k + \mathbf{b} \tag{2}$$

where $\mathbf{z}_k = [o_1, o_2, o_3, o_1 * o_2]^T$ and $\mathbf{M}$ is 3x4. The sample positions in the table refer to the samples used in the correction and are indexed from left to right and then top to bottom.

The result of the enhancement of the faded photo is shown in Figure 8. In this image, the main changes were in the color of the red car, the green grass and leaves. By specifying a small area on the hand, we were able to get some flesh tone in the enhancement.
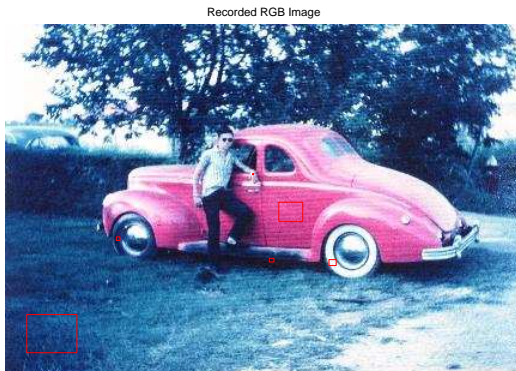
5. REFERENCES

[1] P. Laihanen, "A New Approach to the Manipulation of Colour Display Images,", *SPIE Proceedings*, Vol. 1909, pp. 31-43, Feb. 1993.
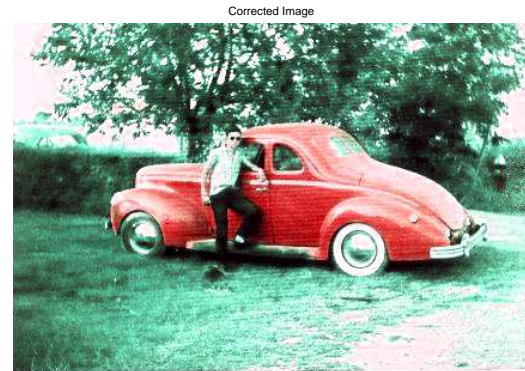
Table 1. $\Delta E94$ errors for correcting Figure 6

Method	Number of Samples	Sample Positions	$\Delta E94$ Error
Affine	4	6 9 41 47	3.5
Affine + RG	5	6 9 35 41 47	3.4
Affine + RG+GB	6	2 6 9 35 41 47	2.6
Affine + RG+GB+RB	7	2 6 9 14 35 41 47	2.4

Fig. 6. Desaturated Munsell image (uniform illumination)

Corrected Image

Fig. 8. Enhanced Faded Photograph

[2] M. J. Vrhel and H. J. Trussell, "Color Device Calibration: A Mathematical Formulation," *IEEE Transactions on Image Processing*, Vol. 8, No. 12, pp. 1796-1806, Dec. 1999.

[3] Vora, P. L., Harville, M. L., Farrell, J. E., Tietz, J. D., and Brainard, D. H. (1997). Image capture: synthesis of sensor responses from multispectral images. Proceedings of the 1997 IS&T/SPIE Conference on Electronic Imaging, (San Jose, CA, February 10-14, 1997), 3018, 2-11. Image is found at http://color.psych.upenn.edu / hyperspectral / bearfruitgray / bearfruitgray.html

Recorded RGB Image

Fig. 7. Original Faded Photograph