Color Correcting Uncalibrated Digital Images

Vlad C. Cardei† and Brian Funt▲

School of Computing Science, Simon Fraser University, Burnaby, B.C., Canada

Color images often must be color balanced to remove unwanted color casts. Color balancing uncalibrated images (e.g. downloaded from the Internet or scanned from an unknown film) adds additional challenges to the already difficult problem of color correction because neither the pre-processing to which the image was subjected, nor the camera sensors or camera balance are known. In this article, we propose a framework for dealing with some aspects of this type of image. In particular, we discuss the issue of color correcting images where an unknown 'gamma' non-linearity may be present. We show that the diagonal model, used for color correcting linear images, also works in the case of gamma corrected images. We also discuss the influence that unknown camera balance and unknown sensors have on color constancy algorithms. To perform color correction on uncalibrated images, we extend previous work on using a neural network for illumination, or white-point, estimation from the case of calibrated images to that of uncalibrated images of unknown origin. The results show that the chromaticity of the ambient illumination in uncalibrated images can be estimated with an average CIE Lab error around 5ΔE. Comparisons are made to the grayworld and white-patch methods.

Journal of Imaging Science and Technology 44: 288–294 (2000)

Introduction

The color of a surface appearing in an image is determined in part by its surface reflectance and in part by the spectral power distribution of the light illuminating it. Thus, as is well known, a variation in the scene illumination changes the color of the surface as it appears in an image. This creates problems for computer vision systems, such as color-based object recognition[1,2] and digital cameras. For a human observer, however, the perceived color shifts due to changes in illumination are relatively small. In other words, humans exhibit a relatively high degree of color constancy.[3]

From a computational perspective, we define the goal of color constancy to be the computation of an image with the same colors (RGB pixel values) as would have been obtained by the same camera for the same scene under a standard 'canonical' illuminant. We see this as a two-stage process: (1) estimate the chromaticity of the illumination; (2) correct the image colors based on this estimate.

Illumination estimation in this sense is also commonly referred to as white-point estimation. Even when the imaging device's characteristics are fully known, accurate illumination estimation for color cast removal has proven difficult, but there has been progress.[4–10] One simple, but often impractical way to estimate the illumination is to have a white patch in the image; the chromaticity of the patch as seen in the image will then be the chromaticity of the illuminant. Other more sophisticated color constancy methods will be discussed below.

Invited paper; Original manuscript received November 10, 1999

Color Plates 5 is printed in the Color Plate Section of this issue, on page 378.

† Student member; ▲ IS&T Member

©2000, IS&T—The Society for Imaging Science and Technology

After estimating the actual illuminant's chromaticity, the image colors can be 'color corrected', that is adjusted to be approximately what they would have been under the canonical illuminant, as shown in **Color Plate 5 (p. 378)**. Color correction is based on a diagonal, or coefficient-rule, transformation. The coefficients of the transformation are computed by comparing the chromaticities of the estimated actual and canonical illuminants.

Color constancy is an under-determined problem and is thus impossible to solve in the most general case. In general, existing color constancy algorithms[4–10] rely in one way or another on accurate camera calibration, as well as on assumptions about the statistical properties of the expected illuminants and surface reflectances. In the case of digital photography images, the camera can be calibrated so that the sensor sensitivities as a function of wavelength are known. As well the sensor response as a function of intensity can be determined. However, in many situations, full calibration is not possible. For example, with images downloaded over the Internet or scanned from film, the imaging characteristics are either unknown or else, as in the case of film, very difficult to control.

Estimating the chromaticity of the illumination in an image of unknown origin poses additional challenges. First of all, not knowing the sensor sensitivity curves of the imaging device means that even for a known surface under a known illuminant, we will not be able to predict its RGB value. Figure 1 shows how much the chromaticities in the rg-chromaticity space (r=R/[R+G+B] and g=G/[R+G+B]) can vary between two cameras. It shows the chromaticities of the Macbeth Colorchecker® patches that would be obtained by a SONY DXC-930 and a Kodak DCS460 camera, both color balanced for the same illuminant. The data for Fig. 1 were synthesized from the known camera sensitivity curves to avoid the values being disrupted by noise or other artifacts.[11,12] Although the white values coincide–as they must, given that cameras were balanced identically–there is a substantial chromaticity difference between the chromaticities from the two cameras for many of the other patches.

A further problem for color constancy on uncalibrated images, is that we do not know the illuminant for which the camera was balanced. Even if two images are taken with the same camera, the output will be different for different color balance settings.

Yet another complication is the question of camera 'gamma'. Cameras whose images are expected to be viewed on a CRT monitor, generally provide output values which are non-linear with respect to the input intensity. The camera non-linearity offsets the CRT non-linearity so that in combination the two non-linearities generate a screen luminance directly proportional to the scene luminance. We must understand the effect this non-linearity has on color correction if we are to handle images of unknown origin.

Of course, in principle, the camera response as a function of intensity could be arbitrary so some assumptions must be introduced. Certainly we can expect the response to be monotonic. We will restrict our attention to "gamma" functions of the form:

$$I = S \cdot D^{\gamma} \tag{1}$$

where I is the resulting pixel value, S is the camera gain, and D is the light intensity value, scaled in the 0.1 range. A typical value of γ is 0.45, however, the results below apply for any value of γ.

Poynton[13] discusses gamma in imaging systems in detail. In what follows, we assume that any non-linearity in the sensor response has been created by a function of the form given in Eq. 1, but that the precise value of the parameter γ is unknown because γ may differ between imaging systems. We will term images for which gamma does not equal unity to be 'gamma-on' images. Linear images are 'gamma-off' and have $\gamma = 1$.

Changing γ shifts most chromaticities, with the general effect being desaturation. Usually $r \neq r^{gamma}$ and $g \neq g^{gamma}$ after the transformation:

$$\begin{cases} r = R/(R+G+B) \\ g = G/(R+G+B) \end{cases} \overset{gamma}{\Rightarrow} \begin{cases} r^{gamma} = R^{\gamma}/(R^{\gamma}+G^{\gamma}+B^{\gamma}) \\ g^{gamma} = G^{\gamma}/(R^{\gamma}+G^{\gamma}+B^{\gamma}) \end{cases} \tag{2}$$

Grays of the form R=G=B, and completely saturated colors such as R=.5, G=B=0, of course remain unchanged by Eq. 2.

The Effect of γ on Color Correction

In terms of the effect of γ on color correction, a crucial question is whether or not the diagonal model, which has been shown to work well on linear image data,[14] still holds once the non-linearity of γ is introduced? We address this question both empirically and theoretically.

Consider an n-by-3 matrix, Q_1, of RGBs from an image of a scene illuminated by E_1, and a similar matrix, Q_2, containing RGBs from the corresponding image of the same scene but illuminated by E_2. In the case where the diagonal model of illumination change holds exactly (see Refs. 14 and 15 for a discussion of the conditions in which this is the case) then there exists a diagonal matrix M such that

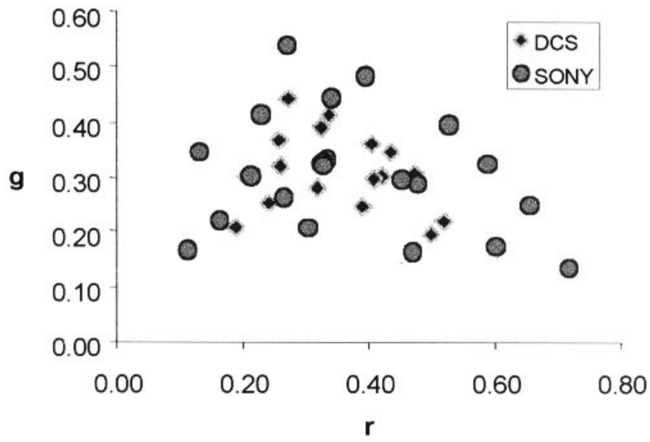

Figure 1. Variation in chromaticity response of two digital cameras.

$$Q_1 \cdot M = Q_2 \tag{3}$$

M depends only on illuminants E_1 and E_2 and does not depend on the RGB values in the images. The diagonal terms of M can be computed by comparing the RGB of any reflectance under the two illuminants for which {R,G,B}> 0. In particular, if $(R_1, G_1, B_1)_{wh}$ is the value of a white reflector under illuminant E_1 and $(R_2, G_2, B_2)_{wh}$ is the corresponding value of the same white reflector under illuminant E_2, then M is given by

$$M = \begin{bmatrix} R_2/R_1 & 0 & 0 \\ 0 & G_2/G_1 & 0 \\ 0 & 0 & B_2/B_1 \end{bmatrix} \tag{4}$$

Let M^{γ} denote element-by-element exponentiation of the elements of matrix M. When the diagonal model M holds exactly for a linear image ($\gamma = 1$), then it will also hold exactly for the same image after a transformation of the form of Eq. 1 with $\gamma \neq 1$. In this case, the diagonal transformation matrix becomes M^{γ}:

$$Q_1^{\gamma} \cdot M^{\gamma} = Q_2^{\gamma} \tag{5}$$

This equality becomes obvious if we expand it:

$$\begin{bmatrix} R_{11}^{\gamma} & G_{11}^{\gamma} & B_{11}^{\gamma} \\ R_{12}^{\gamma} & G_{12}^{\gamma} & B_{12}^{\gamma} \\ \vdots & \vdots & \vdots \end{bmatrix} \cdot \begin{bmatrix} \lambda_R^{\gamma} & 0 & 0 \\ 0 & \lambda_G^{\gamma} & 0 \\ 0 & 0 & \lambda_B^{\gamma} \end{bmatrix} = \begin{bmatrix} R_{21}^{\gamma} & G_{21}^{\gamma} & B_{21}^{\gamma} \\ R_{22}^{\gamma} & G_{22}^{\gamma} & B_{22}^{\gamma} \\ \vdots & \vdots & \vdots \end{bmatrix} \tag{6}$$

which is true, because:

$$R_{1k}^{\gamma} \cdot \lambda_R^{\gamma} = R_{2k}^{\gamma} \text{ and } G_{1k}^{\gamma} \cdot \lambda_G^{\gamma} = G_{2k}^{\gamma} \text{ and } B_{1k}^{\gamma} \cdot \lambda_B^{\gamma} = B_{2k}^{\gamma} \tag{7}$$

where k is a row index in matrices Q_1 and Q_2 and λ is an element of the diagonal matrix M.

In practice, the diagonal model is not perfect, so the transformation matrix will also contain small off-diagonal terms.[14,15] Introducing γ amplifies these off-diagonal terms. To explore the effects of γ on the off-diagonal terms, we synthesized two images under different illuminants and evaluated the accuracy of the diagonal transformation in mapping between them. The two images were generated using measured spectral reflectances of the 24 patches of the *Macbeth Colorchecker*. One image was synthesized relative to CIE illuminant A and the other relative to D65, both using the spectral sensitivities of the SONY DXC-930 camera scaled so the resulting RGBs fall in [0...1].

If A is the matrix of synthesized RGBs under illuminant A and D is the matrix of corresponding RGBs under illuminant D65, the transformation from matrix D to A is given by

$$D \cdot M = A \tag{8}$$

For linear image data, the best (non-diagonal) transformation matrix M and the best diagonal matrix M_D (in the least square errors sense) are found to be:

$$M = \begin{bmatrix} 4.225 & 0.166 & -0.076 \\ -0.372 & 2.027 & 0.132 \\ 0.045 & -0.048 & 0.792 \end{bmatrix}$$

and

$$M_D = \begin{bmatrix} 3.886 & 0 & 0 \\ 0 & 2.036 & 0 \\ 0 & 0 & 0.821 \end{bmatrix}. \tag{9}$$

These transformation matrices are computed to minimize the mean square error using the pseudo-inverse:

$$M = D^* \cdot A, \tag{10}$$

where "*" denotes the pseudo-inverse of the matrix.

The accuracy of the matrix transformation is measured by comparing the estimated RGBs, $E = DM$, with A, the actual RGBs under illuminant A. For the non-diagonal case, the average error $\mu_{\text{linear}} = 0.0088$ and the standard deviation $\sigma_{\text{linear}} = 0.0061$. In the perceptually uniform CIE Lab space the average ΔELab error obtained from the same matrix transformation is $\mu_{\text{Lab}} = 2.14$, with standard deviation $\sigma_{\text{Lab}} = 1.56$.

In Eq. 9, the diagonal elements of M_D resemble those of M, but do not equal them. The differences compensate for the effect of constraining the non-diagonal terms to 0. Of course, when we are restricted to a diagonal matrix we expect the error to be somewhat larger than when using a non-diagonal matrix. Using M_D, the average error $\mu'_{\text{linear}}=0.0192$ and the standard deviation $\sigma'_{\text{linear}} = 0.0128$. In CIE Lab space the average error $\mu'_{\text{Lab}}=3.36$ and the standard deviation $\sigma'_{\text{Lab}}=2.30$. Although these errors are almost twice as large as for the full non-diagonal linear transformation, they are still quite small and show that a diagonal transformation provides a good method of predicting the effects of an illumination change on an image.

To determine the effect of γ on the effectiveness of the diagonal model, we took the previously synthesized data and applied γ of 1/2.2. In this case the best transformation M_γ and the best diagonal transformation $M_{D\gamma}$ are

$$M_\gamma = \begin{bmatrix} 2.020 & 0.086 & -0.043 \\ -0.204 & 1.381 & 0.095 \\ 0.038 & -0.052 & 0.877 \end{bmatrix}$$

and

$$M_{D\gamma} = \begin{bmatrix} 1.855 & 0 & 0 \\ 0 & 1.380 & 0 \\ 0 & 0 & 0.914 \end{bmatrix} \tag{11}$$

The average error in RGB using the full 3-by-3 transformation, $M\gamma$, is $\mu_{\text{gamma}} = 0.0067$ with standard deviation $\sigma_{\text{gamma}} = 0.0037$. In CIE Lab the average ΔELab error is $\mu_{\gamma\text{Lab}} = 1.06$ with standard deviation $\sigma_{\gamma\text{Lab}} = 0.69$. The average error in RGB when restricted to the diagonal transformation, $M_{D\gamma}$ becomes $\mu'_{\text{gamma}} = 0.0180$ with standard deviation $\sigma'_{\text{gamma}} = 0.0103$. In CIE Lab space the average ΔELab error is $\mu_{\gamma\text{Lab}} = 2.04$ with standard deviation $\sigma_{\gamma\text{Lab}} = 1.39$. The magnitude of these errors for the gamma-on case is comparable to those of the linear, gamma-off case above.

The above results are summarized in the charts of Figs. 2 and 3. From these charts it is clear that the diagonal model still holds for images to which a non-linear γ has been applied even where the diagonal model in the linear case provides only an approximate model of illumination change. Transformation errors for non-linear images are smaller than for linear ones.

Another issue that must be considered in terms of color correcting images of unknown γ has to do with the effect that a scaling in brightness of the form (R,G,B) to (kR,kG,kB) might have. A brightness scaling may result from a change in incident illumination or a change in camera exposure settings. Also the image intensities may have been scaled by a user simply to make it look better.

Whatever the cause, it turns out that a brightness change does not affect a pixel's chromaticity even in the case of gamma-on images. Consider a pixel (R,G,B) from a linear image with red chromaticity of r = R/(R + G + B). After γ, its red chromaticity will be:

$$r^{gamma} = R^\gamma \Big/ \left(R^\gamma + G^\gamma + B^\gamma\right) \tag{12}$$

In the linear case, any brightness scaling leaves the chromaticity unchanged. In the non-linear case, the red chromaticity of the pixel will be:

$$\begin{aligned} r_N^{gamma} &= (kR)^\gamma / ((kR)^\gamma + (kG)^\gamma + (kB)^\gamma) \\ &= R^\gamma / (R^\gamma + G^\gamma + B^\gamma) = r^{gamma} \end{aligned} \tag{13}$$

Similar results hold for other chromaticity channels, so brightness changes do not effect the chromaticities in gamma-on images. Note, however, that this does not mean that the chromaticity of a pixel is the same before and after the application of γ.

Color Correction on Non-Linear Images

We have shown thus far that for both linear and non-linear image data the diagonal model approximation holds and that changing intensity does not affect chromaticity in either case. In what follows, we will address the commutativity of γ and color correction. Given an image I, represented as an n-by-3 matrix of RGBs, we define two operators on this image. $\Gamma(I)$ denotes the application of γ and $C(I,M)$ denotes the color correction operator:

$$\Gamma(I) = I^\gamma \text{ and } C(I,M) = I \bullet M. \tag{14}$$

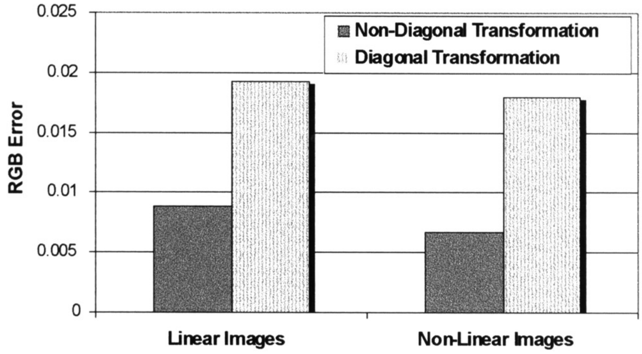

Figure 2. Error in predicting the effects of illumination change on image data in RGB space for both linear and non-linear image data comparing general 3-by-3 matrices to diagonal ones.

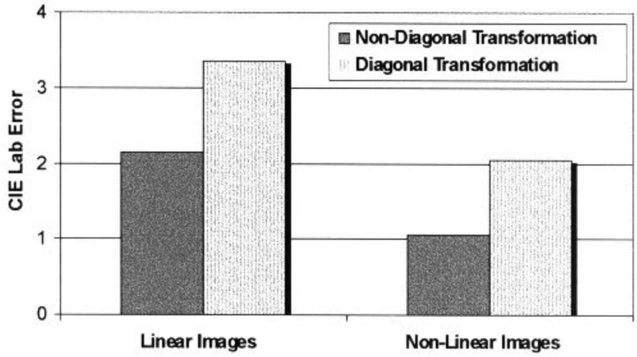

Figure 3. CIELAB ΔELab error in predicting the effects of illumination change on image data for both linear and non-linear image data comparing general 3-by-3 matrices to diagonal ones.

We wish to find out if the two operators commute, i.e. if

$$C\big(\Gamma(I), M\big) = \Gamma\big(C(I, M)\big). \tag{15}$$

The diagonal transformation matrix M depends on the image I and the illuminant under which it was taken. This transformation maps pixels belonging to a white surface in the image into achromatic RGB pixels *(N,N,N)*.

If (R_{wh}, G_{wh}, B_{wh}) is the color of the illuminant (i.e., the camera's response to an ideal white surface under that illuminant) for image I and (R, G, B) is an arbitrary pixel in I, then

$$\begin{aligned} C\big(\Gamma([R,G,B]), M_\gamma\big) &= C\big([R^\gamma, G^\gamma, B^\gamma], M_\gamma\big) \\ &= \big[m_\gamma^R R^\gamma, m_\gamma^G G^\gamma, m_\gamma^B B^\gamma\big] \end{aligned}, \tag{16}$$

where

$$M_\gamma = \begin{bmatrix} m_\gamma^R & 0 & 0 \\ 0 & m_\gamma^G & 0 \\ 0 & 0 & m_\gamma^B \end{bmatrix} \tag{17}$$

is the transformation to be used on images with gamma-on.

If we know the color of the illuminant, the diagonal elements of M_γ can be computed from the following equation:

$$\begin{aligned} C\big(\Gamma([R_{wh}, G_{wh}, B_{wh}]), M_\gamma\big) &= C\big([R_{wh}{}^\gamma, G_{wh}{}^\gamma, B_{wh}{}^\gamma], M_\gamma\big) \\ &= \big[m_\gamma^R R_{wh}{}^\gamma, m_\gamma^G G_{wh}{}^\gamma, m_\gamma^B B_{wh}{}^\gamma\big] = [1,1,1]. \end{aligned} \tag{18}$$

Thus, the transformation matrix becomes:

$$M_\gamma = \begin{bmatrix} 1/R_{wh}^\gamma & 0 & 0 \\ 0 & 1/G_{wh}^\gamma & 0 \\ 0 & 0 & 1/B_{wh}^\gamma \end{bmatrix}. \tag{19}$$

We can rewrite Eq. 16, as a function of *(R,G,B)* and (R_{wh}, G_{wh}, B_{wh}):

$$\begin{aligned} C\big(\Gamma([R,G,B]), M_\gamma\big) &= \big[m_\gamma^R R^\gamma, m_\gamma^G G^\gamma, m_\gamma^B B^\gamma\big] \\ &= \left[\frac{1}{R_{wh}^\gamma}\cdot R^\gamma, \frac{1}{G_{wh}^\gamma}\cdot G^\gamma, \frac{1}{B_{wh}^\gamma}\cdot B^\gamma\right]. \end{aligned} \tag{20}$$

The right hand side of Eq. 15 can be written as:

$$\Gamma\big(C(I, M)\big) = \Gamma\big([m^R R, m^G G, m^B B]\big), \tag{21}$$

where m^x are the diagonal elements of matrix M.

Because M maps a white surface into white, we can write M as:

$$M = \begin{bmatrix} 1/R_{wh} & 0 & 0 \\ 0 & 1/G_{wh} & 0 \\ 0 & 0 & 1/B_{wh} \end{bmatrix} \tag{22}$$

Thus, Eq. 21 can be rewritten as:

$$\begin{aligned} \Gamma\big(C([R,G,B], M)\big) &= \Gamma\left(\left[\frac{1}{R_{wh}}R, \frac{1}{G_{wh}}G, \frac{1}{B_{wh}}B\right]\right) \\ &= \left[\frac{1}{R_{wh}^\gamma}\cdot R^\gamma, \frac{1}{G_{wh}^\gamma}\cdot G^\gamma, \frac{1}{B_{wh}^\gamma}\cdot B^\gamma\right] \end{aligned} \tag{23}$$

From Eqs. 20 and 23 it follows that Eq. 15 is true for any pixel in I, i.e., that color correction and γ application are commutative. As a result, we can color correct gamma-on images in the same way as linear images.

In the equations above we assumed that there is a perfect white surface in the image I or, equivalently, that the color of the illuminant is known. Given knowledge of the illuminant, the color correction method is the same for gamma-on and gamma-off images, although the parameters of the transformation differ. Color constancy methods used to determine the illumination, however, must change, because γ affects a pixel's chromaticity. In general, applying γ results in a more desaturated color.

The difference in chromaticities throughout an image will result in a different statistical distribution of chromaticities. This change in the distribution of chromaticities can adversely affect those color constancy algorithms that rely on *a priori* knowledge about the statistics of the world.[5–9]

Color Correcting Images from Unknown Sensors

There are two aspects related to unknown sensors: the color balance of the camera and the sensor sensitivity curves. In most cases, the color balance is determined by scaling the three color channels according to some predetermined settings. The goal of color balance is to obtain equal RGB values for a white patch under a canonical light. In this case, we say that the camera is calibrated for that particular illuminant. Color correcting images taken with an unknown balance does not pose a problem, because the calibrating coefficients can be absorbed in the diagonal transformation that performs the color correction. However, finding the diagonal transformation might prove difficult for stochastic algorithms[5–9] that can have difficulties in generalizing their estimates if they fall outside the illumination gamut for which they were trained.

If the spectral sensitivities of the camera sensors are unknown, many color constancy algorithms will have difficulty providing reasonable estimates of the scene illumination. As described in the next section, we tested several algorithms on uncalibrated image data and found that the neural network approach works quite well.

Illumination-Estimation Algorithms

We test several different illumination-estimation algorithms on a database of 'uncalibrated' images. The images are uncalibrated in the sense that the imaging characteristics are not provided to the algorithms, even though we have the calibration parameters available so that we can evaluate the results. In particular, we test the white patch algorithm (WP), a version of the grayworld algorithm (GW) and two neural-network-based methods. The gamut-constraint methods[6,7] were not tested because they require information about the expected gamuts of reflectances or illuminants. This information can not be obtained without knowing the sensor sensitivity functions of the devices that acquired the images.

The image database contains 116 images taken with a Kodak DCS-460 camera and 67 images scanned with a Polaroid Sprintscan 35+ slide scanner from various film types: Kodak Gold, Kodak Royal, Agfa Optima, Polaroid HiDef and Fuji Superia. The photographs were scanned using a 'generic' pre-defined scanner setting. This setting is consistent with the assumption of unknown pre-processing. Using the manufacturer's optimal setting for each specific film type would have allowed the scanner driver to accommodate partially for the differences in film.

We divide the image database into two sets, the first for training and the second for testing. The training set contains 102 images and is used for training the neural network and computing the average color for use in the database grayworld algorithm. The test set contains the other 81 images (57 DCS images and 24 slides).

For the GW algorithm, the chromaticity of the illuminant is determined from the average of all the pixels in an image. GW assumes that the average color of the scene is gray and that any departure from this average in the image is caused by the color of the illuminant. The average is computed relative to the average chromaticity computed using all pixels in the training database. Using the database average as the definition of gray compensates for the fact that gray may not have exactly equal r and g chromaticities. Nonetheless, GW's performance will be poor when a test image has a different average distribution than the images used for computing the database average.

The WP algorithm determines white, and hence the illuminant color, as the maximum R, maximum G and maximum B found in the image. The WP algorithm has roots in the family of retinex algorithms,[4] but it is only equivalent to it under restricted circumstances.

Two differently trained neural networks were used for illumination estimation. The network architecture was the same in both cases; namely, a Perceptron with two hidden layers as we have previously described.[8,11,17] The networks are trained to estimate the chromaticity of the illuminant based on the binarized rg-chromaticity histogram of an input image. The 3600-node input layer is fed binary values representing the presence or absence of chromaticities falling within a particular chromaticity bin. The first hidden layer contains 50 neurons and the second layer 20 neurons. The output layer consists of only two neurons representing the chromaticity of the illuminant. All neurons have a sigmoid activation function.

Both neural networks were trained using back-propagation. The error function for training and testing is the Euclidean distance in rg-chromaticity space between the actual illuminant and its estimate.

The difference between the training of the two networks concerns the method of determining the actual illuminant. For the first network, the illuminant chromaticity is simply measured from the reference white standard that was placed within each image. This reference white was manually removed from the images before testing the color constancy algorithms. It provides an accurate value for the illuminant's chromaticity. For the second network, a less accurate method is used, which we have called the bootstrapping method.[17] The bootstrapped network uses the GW algorithm to "measure" the chromaticity of the illuminant for training. Clearly, the illuminant value determined by GW will only be approximately correct; nonetheless, previous experiments with calibrated image data showed that the network "learned" to make a better estimate than the simple GW algorithm used to train it. Our new experiments described below show that bootstrapping works even for the more general case of non-linear images acquired from various sources. This approach allows us to train a neural network for a range of uncalibrated cameras and scanners, without having to explicitly measure white patches in the set of training images.

Experimental Results

The algorithms presented above were tested on an image database containing 81 images. Figures 4 and 5 show the relative performance of the color constancy algorithms. The figures show the average errors over the whole test set as well as for each type of input (i.e. for DCS images and slides). In Fig. 4, the average errors are computed in the rg-chromaticity space, the same space in which the neural network was trained. "Nothing" refers to the assumption that the illuminant is the one for which the device is calibrated and reflects the variation in the chromaticity of the illuminant across the test set of images, relative to

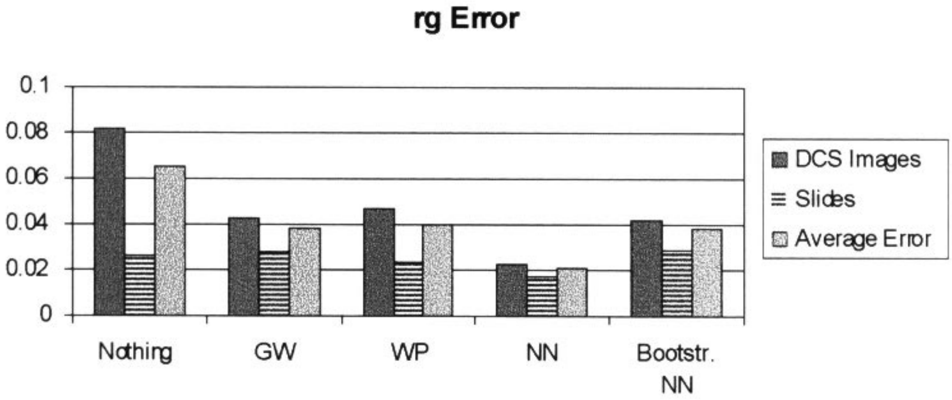

Figure 4. Average errors measured in rg-chromaticity space.

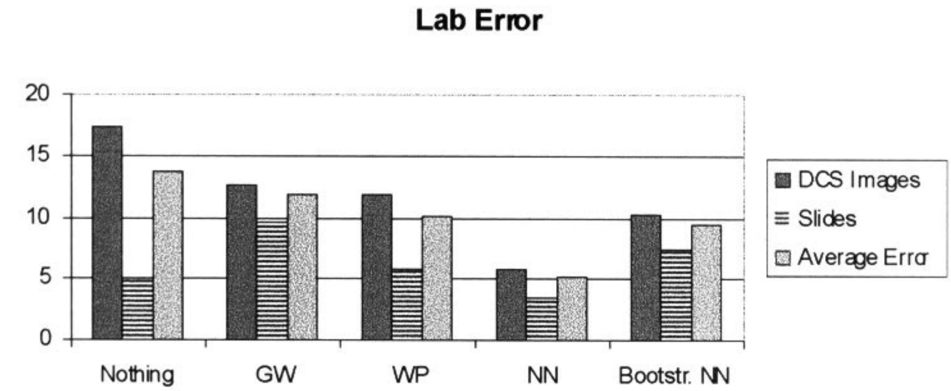

Figure 5. Average CIE Lab ΔE space between actual and estimated illuminant fixed to the same L* value.

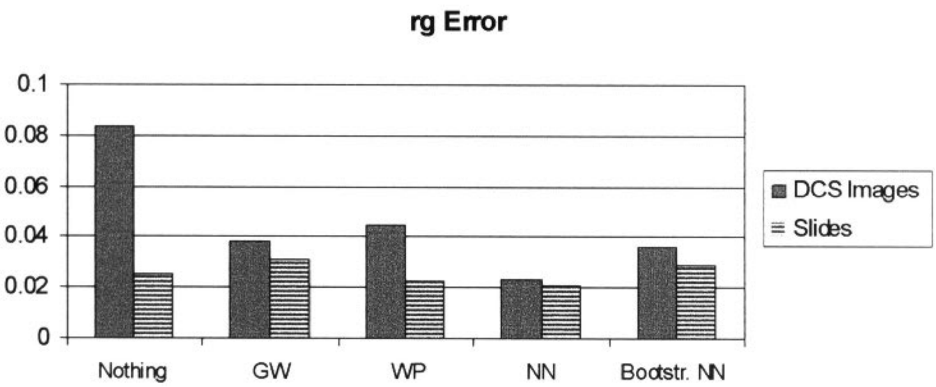

Figure 6. Average error in rg space when the training and test data come from the same uncalibrated source.

white (located at r=g=1/3 in rg-chromaticity space). "NN" refers to the neural network trained with accurately measured illumination data, while "Bootstrapped NN" refers to the same network trained using GW illumination estimates.

Figure 5 presents similar results, but with the error measured in CIE Lab space. The conversion from the RGB space to CIE Lab assumes the images are to be viewed on a sRGB-compliant[16] monitor.

Figures 6 and 7 compare the results of neural networks trained on images from a single uncalibrated device (i.e., camera or scanner) with the other algorithms. The results show that in this case, the accuracy of the neural network is much better than when the device type varies.

Discussion

We presented a framework for dealing with a quite general case of color correction; namely, that of images for which both the spectral sensitivity of the sensors and γ setting are unknown. One conclusion is that for images to which γ has been applied, it is possible to perform color correction by a diagonal transformation without

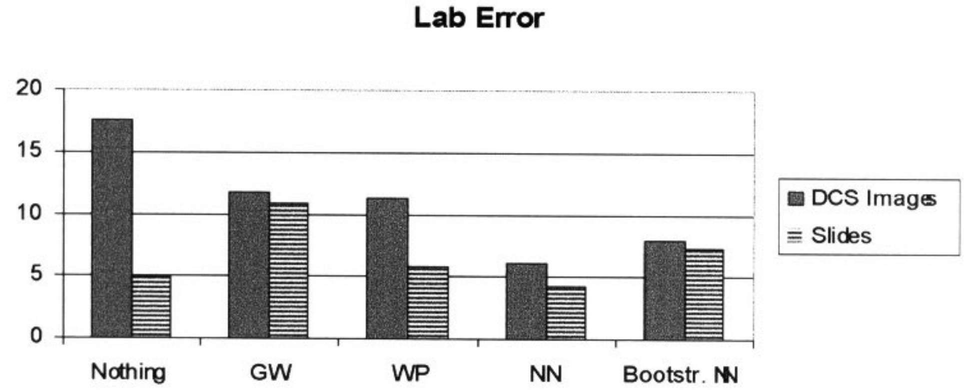

Figure 7. Average CIE Lab ΔE space between actual and estimated illuminant fixed to the same L* value when the training and test data come from the same uncalibrated source.

first linearizing the image data. The off-diagonal elements of the general image transformation are larger with respect to the diagonal elements when γ has been applied and thus the average error of a diagonal transformation (which ignores the off-diagonal terms) will increase. However, the perceptual error is still very small and the diagonal transformation thus remains a good model of illumination change.

In the case of unknown sensors, as we saw in Fig. 1 there are large differences in sensor response, even for cameras balanced for the same illuminant. This variation in the distribution of sensor responses can adversely affect color constancy algorithms that rely on assumed distributions of sensor responses.

Previous studies[8,11,17] based on calibrated, linear image data have shown that a neural network can accurately estimate the illumination chromaticity. Often we must work with uncalibrated image data, so we trained and tested several algorithms on uncalibrated data, but in a controlled manner. On this test data, the neural net average error is 5.14ΔELab. We believe this to be useful for removing color casts from images of unknown origin. In the tests with the bootstrapping method of training the neural network, the ΔELab error increased to 9.38. Nonetheless, this is better than either the GW or WP methods and the bootstrapping method can be applied in situations where accurate measurements of the illuminant chromaticity are unavailable for training.

Acknowledgment. The authors gratefully acknowledge the financial support of the Natural Sciences and Engineering Research Council of Canada and Hewlett Packard Incorporated, and Lindsay Martin and Fred Kyba for their help acquiring much of the image data used in the experiments.

References

1. M. Swain and D. Ballard, Color Indexing, *Int. J. of Computer Vision*, **7**(1), 11–32 (1991).
2. B. Funt, K. Barnard and L. Martin, Is Colour Constancy Good Enough?, H. Burkhardt and B. Neumann (eds.), *Proc. ECCV '98, Lecture Notes in Computer Science*; Vol. 1406, Springer-Verlag, 1998, pp. 445–459.
3. D. H. Brainard, W. A. Brunt and J. M.. Speigle, Color Constancy In The Nearly Natural Image. 1. Asymmetric Matches, *J. Opt. Soc. Am. A*, **14**(9), 2091–2110 (1997).
4. E. H. Land, The Retinex Theory of Color Vision, *Scientific American*, 108–129 (1977).
5. G. Buchsbaum, A Spatial Processor Model for Object Colour Perception, *J. Franklin Inst.* **310** (1), 1–26 (1980).
6. D.A. Forsyth, A Novel Algorithm for Color Constancy, *Int. J. of Computer Vision,* **5** (1), 5–36 (1990).
7. G. Finlayson, Color in Perspective, *IEEE Trans. PAMI*, **18**(10), 1034–1038 (1996).
8. B. Funt, V. Cardei, and K. Barnard, Learning Color Constancy, *Proc. Fourth Color Imaging Conf.*, IS&T, Springfield, VA, 1996, pp. 58–60.
9. G. Finlayson, P. Hubel and S. Hordley, Color by Correlation, *Proc. Fifth Color Imaging Conf.*, IS&T, Springfield, VA, 1997, pp. 6–9.
10. V. Cardei, B. Funt and K. Barnard, White Point Estimation for Uncalibrated Images, *Proc. Seventh Color Imaging Conf.*, IS&T, Springfield, VA, 1999, pp. 97–100.
11. V. Cardei, B. Funt and K. Barnard, Adaptive Illuminant Estimation Using Neural Networks, L. Niklasson, M. Bodén and T. Ziemke (eds.), *Perspectives in Neural Computing; Vol. 2, Proc. of the 8th Int. Conf. on Artificial Neural Networks*, Springer-Verlag, London, 1998, pp. 749–754.
12. B. Funt, V. C. Cardei and K. Barnard, Neural Network Color Constancy and Specular Reflecting Surfaces, *AIC Color '97*, 523-526 (1997).
13. C. Poynton, The Rehabilitation of Gamma, *Proc. SPIE* **3299**, 232–249 (1998).
14. G. Finlayson, M. Drew, and B. Funt, Color Constancy: Generalized Diagonal Transforms Suffice, *J. Opt. Soc. Am. A*, **11**(11), 3011–3020 (1994).
15. J. A. Worthey and M. H. Brill, Heuristic Analysis Of Von Kries Color Constancy, *J. Opt. Soc. Am. A*, **3**(10), 1709–1712 (1986).
16. M. Anderson, R. Motta, S. Chandrasekar and M. Stokes, Proposal for a Standard Default Color Space for the Internet—sRGB, *Proc. Fourth Color Imaging Conf.*, IS&T, Springfield, VA, 1996, pp. 238–246.
17. B. Funt and V. Cardei, Bootstrapping Color Constancy, *Proc. SPIE,* **3644,** 421–428 (1999).

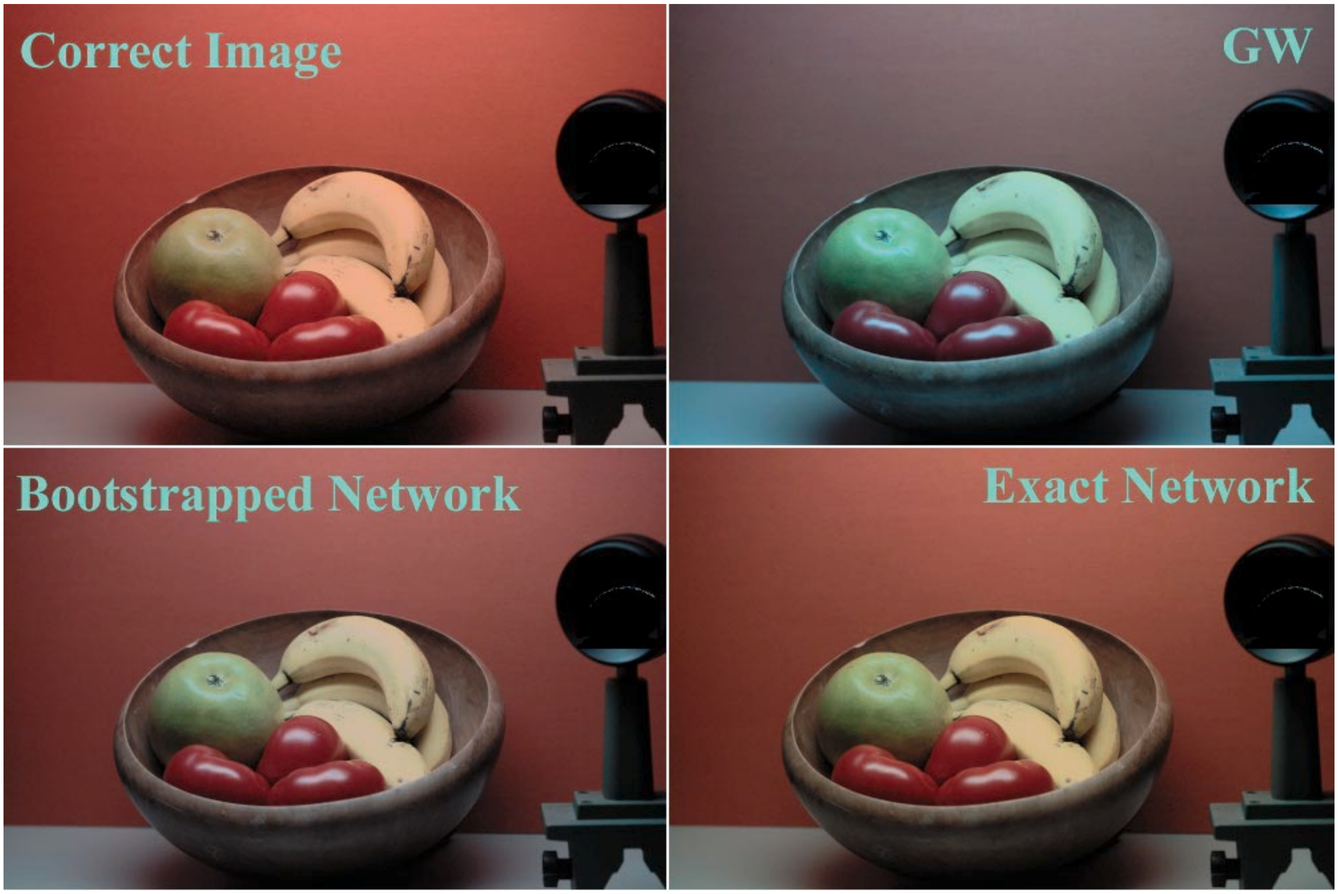

Plate 5. Example of color correction based on various methods including grey world, a neural network trained on estimated training data via bootstrapping,[17] and a neural network trained on accurate training data.[8,11] (Cardei and Funt, pp. 288–294).