Colour Correction for an Image Sequence

Glen Pringle
Dept. of Computer Science, Monash University, Australia

Binh Pham
School of Computing and Information Technology, Griffith University, Australia

Abstract

A technique for correcting the colours in an image sequence to improve their visual appearance when viewed together is presented. The technique constructs a mapping from the registered overlap regions of two adjacent images using least-squares regression to minimize the colour differences. Strategies for compensating for colours not present in the overlapping region and techniques for reducing noise and errors caused by values outside the colour gamuts are also described.

1. Introduction

When image sequences in the spatial domain are viewed together, they often appear unnatural due to colour differences between the images. For example, panoramic photos, when joined together, may have large changes in the colours between individual photos due to factors such as changes in the illumination (e.g. shadows, clouds) or settings on the camera (e.g. shutter speed, aperture). Colour differences may also occur when several steps are involved in the capture of the images before they are in a usable form in the computer (e.g. film processing, printing, scanning). At each of these steps, the colour may be balanced to enhance the appearance of each individual image without taking into account the overall result of the sequence. To minimize the discrepancy in visual appearance of adjacent images, the colours need to be corrected.

Most research work done on colour correction up-to-date is related to the problem of device-to-device mapping. This usually involves finding a mapping to match a set of specified output colours to that of known input colours. This mapping will then be used for correcting other colours. Matrix methods are commonly employed for this task, and the matrix values are obtained from least squares techniques.

Kang[1] presented a method of scanner calibration and examined the errors associated with the scanner and this process. The scanner was used to obtain RGB values for each of the test patches. These values were grey-balanced, by using a least-squares fit for the grey patches in the test target. In the case of the grey balanced curve being slightly convex, a curve fit was used. The grey-balanced RGB values are used in conjunction with the known colorimetric CIE/XYZ values to obtain a transfer matrix through multiple polynomial regression. The patch values were transformed using this transfer matrix and then further transformed into CIE/LAB space. These values were compared to the measured values to give an average colour difference. It was found that large numbers of test colours did not affect accuracy of the calibration significantly when a well-chosen set was used.

Schwartz[2] carried out the colour correction/equalization for input and output devices based on the use of a small carefully selected set of *touchstones,* which are specific known input colour values. Alternatively, a set of test colours which have an accurately control spectra such as the Macbeth ColorChecker could be used. If CIE XYZ coordinates are used, colours which differ to the touchstones should have good interpolation. If only three touchstones are available, corrections for tristimulus values which are close to these colours will be very accurate but the results may be poor for other values. With a higher number of touchstones, a least-squares method which minimizes the sum of squared errors of the corrected values can be applied. Results obtained for the tristimulus values close to the touchstones were very good, with an almost imperceivable difference between the corrected colours and the desired colours.

Both these methods rely on having a set of test colours with known colorimetric values which are not available in our case. Instead, we use the information present in the overlap regions of two adjacent images in the sequence. Our method for correcting the colour differences in pairs of adjacent images involves two main steps. The first step is to find the relationship between the registered overlap regions of two adjacent images in the sequence using multiple regression. The derived mapping is then used in the second step to correct the colour in one image to match it to its adjacent image. The results can be further improved by using a technique which we termed *incremental regression* which finds the intermediate mappings to match the adjacent image pairs gradually in two or more stages. To compensate for the colours which are only present in the non-overlap regions of images, a set of representatives of such colours can also be used in the process of constructing the mapping. Sigmoid functions can also be employed to transform the data before applying the mapping and in-verse transform is performed on the outcome of the mapped data to obtain the final result. This would reduce the problem caused by colour values which fall outside the colour gamuts.

2. Finding Mapping

In order to find the mapping for correcting the two adjacent images, the images must have a small overlap region which can be found by registering the images (images A' and B' in Figure 1). Once the images are registered, the pixel values at the same location in each image are known. We wish to find a mapping which will convert the colour value of each of these pixels in the overlap region in A' to the corresponding pixel colour value in B'. In some cases, this mapping may be quite simple, e.g. lightening or darkening an image by adjusting the luminance channel. However, in general, this mapping is often rather complex as the causes for the colour differences are compound or unknown. In such cases, the mapping can be modelled using high degree polynomials.

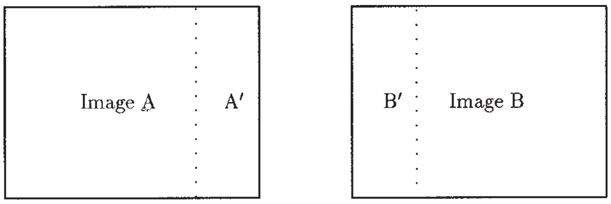

Figure 1. Correction of two adjacent images

Regression analysis[5,6] is a common technique used for finding a relationship between two sets of numbers. Linear regression takes two sets of numbers

$$(x_i), (y_i), {}_i = 1 \ldots n$$

and produces a linear relationship between them, finding the best straight line which approximates the relationship between the data sets by obtaining the set of linear coefficients b_0 and b_l which minimize the sum of squared errors. Thus, if we have observations

$$(x_{i\,,}\, y_i), \ldots (x_{n,}\, y_{n\,)}$$

we want to find β_0 and β_l where

$$y_i = \beta_0 + \beta_{\mathrm{i}}\, x_i + e_i, \tag{1}$$

$i = 1 \ldots n$ (e_i = error term) such that we minimize the sum of squared errors

$$S = \sum_{i=1}^{n} (y_i - \beta_0 - \beta_1 x_i)^2 \tag{2}$$

The estimators b_0 and b_1 of $\beta_{0\ \mathrm{and}}\ \beta_1$ can be easily found:

$$b_0 = \bar{y} - b_1 \bar{x}$$

$$b_1 = \frac{\sum_{i=1}^{n} (x_i - \bar{x}) y_i}{\sum_{i=1}^{n} (x_i - \bar{x})^2} \equiv \frac{\sum_{i=1}^{n} (x_i - \bar{x})(y_i - \bar{y})}{\sum_{i=1}^{n} (x_i - \bar{x})^2}$$

This technique can be extended to find non-linear relationships between multiple sets of data, using a polynomial regression model of the form

$$y = \beta_0 + \beta_1 x + \beta_2 x^{\,2} + \cdots + \beta_n x^{\,n} + e \tag{3}$$

where e is the error term. This equation can be expressed as $y = BX$ where y, B and X are matrices. The data in each row of y is to be matched with that in each column of X. The number of rows of y (or the number of columns of X) is the number of data sets to be matched whereas the number of columns of y (or the number of rows of X) is the number of data items in each data set. The solution to this equation is

$$B = (X^{\,T} X\,)^{-\infty}\, X^{\,T} y \tag{4}$$

with corresponding sum of squared errors,

$$S = (y^T\, y)^{\,-}\, B^T X^{\,T} y \tag{5}$$

This technique, which is called *multiple polynomial regression,* will be used to find the mapping between the pixels in the overlap regions. The coefficients of the polynomial are obtained using Equation 4.

To apply this technique to our problem, the colour components of the pixels in region A' of Figure 1 are used as the values for the matrix X and the colour components of the pixels in region B' for matrix y. Once the matrix B is found, we have the mapping from region A' to B' in the form of Equation 3. This mapping is then applied to every pixel in region A' to correct its colour to produce an image which closely matches region B'.

The choice of the polynomial used (Equation 3) is very important. If it is too simple, it may not be able to adequately model the physical effects of the colour changes. In our experience, a 14-term polynomial with terms up to n^3 provided a good model for most cases.

For special cases where the colour differences in spatial image sequences are mainly caused by the differences in the luminance (e.g. one image being lighter than an adjacent one), performing the regression and mapping in the luminance component of the CIE Lab or Luv spaces should give a more accurate mapping.

We also carried out the colour correction using several different colour spaces: RGB, xyz, CIE Lab and Luv. The results in all these spaces were similar, with the exception that the rgb space did not give as good interpolation for colours which were not present in the overlap region.

3. Refinement Strategies

3.1 Colour Compensation

The mapping B found above can subsequently be applied directly to the whole of Image A for colour correction. However, very often the non-overlap region has different characteristics from the overlap region. It might contain objects which have colours which are not present in the overlap region. Applying the mapping to such colours can produce unexpected effects such as unusual colours or non-existent colours when the values go outside the range of the colour space. For example, the overlap region of two images may not contain red, but in the non-overlap region of one image there is a red object. When the mapping is applied to this red object, the resulting values may go out-of-gamut or produce an unnatural colour. These unexpected effects occur because the mapping found is one which minimizes the sum of squared errors for only the pixel values in the overlap region but does not take into account any pixel values at other positions. Thus, a good representative set of colours is needed for a successful regression scheme.

We need to find an effective way of incorporating colours from the non-overlap region into the mapping. This can be done by adding selective pixel values from the non-overlap region to those in the overlap region when applying regression. These extra pixels which are assigned to map onto themselves are the representatives of the colours which are not present in the overlap region. Thus the resultant mapping will not only correct colours in the overlap region, but should also set the colours we had problems with before to a similar colour. An important task is to find the correct ratio of pixels from the non-overlap region to pixels from the overlap region. If this ratio is too high, it will result in little or no correction, as the best mapping (in terms of sum of squared errors) to convert some colours to themselves is the identity. If the ratio is too low, we would still have the problem of out-of-gamut values.

We have experimented with a compensation factor which takes into account of two quantities: (i) the ratio between the number of pixel values which are unique to the non-overlapping region and the number of pixel values in the overlap region; (ii) the average Euclidean distance between the unique pixels and the pixels in the overlap region. The compensation factor should be high if both the quantities are large and should be low if either or both of these quantities are low.

3.2 Incremental Regression

By its very nature, regression is not iterative since the coefficients are obtained by minimizing the sum of squared errors. However, finding the mapping in one step causes problems when the sample of colours in the overlap region is poor and the amount of colour differences are large. In such situations, a method which allows the correction to be done in stages to cater for a smaller change at each stage, is more desirable. For example, a two-stage method involves an initial mapping to be found which converts initial values to halfway values and then a second mapping which converts halfway values to the final image. This process may involve several stages and consequently, provide several mappings which must be applied sequentially to the image to correct it. Alternatively, since we want to perform the regression on images with smaller differences in order to decrease the chance of having out-of-gamut values, both images can be corrected to a half-way image where they will match. Thus we effectively reduce the amount of correction applied. In some cases, the colours present in the overlap region may provide better representatives for colours in one image, so we may apply more steps to this image and less steps to the adjacent image. This method which we call *incremental regression* was found to reduce noise and the number of out-of-gamut values significantly in our tests.

3.3 Sigmoid Functions

When a value in the non-overlap region does go out-of-gamut, we must replace it with some value which lies within the colour space. This value must be similar to the original value and the corrected surrounding values. For a single pixel which is out-of-gamut, we can use the average of the original and surrounding pixels, but in situations where large sections of the image are out-of-gamut, we must bring the pixels back into gamut. This can be achieved by moving the out-of-gamut values towards the reference white until they are in-gamut.

Another strategy is to use some sigmoid functions to transform the pixel values before regression is performed. These functions have the special property that when they are applied, a set of data which varies slowly at the ends of its range is transformed into another set of data which varies at a much faster rate towards the ends of the range. Hence, there is more scope for variation at the ends of the range when regression is applied on the transformed data. Once the regression is performed, the results can be transformed back to the original range of values, using the inverse sigmoid function. We have experimented with two sigmoid functions:

$$f\ (x) = tan\ (x) \tag{6}$$

and

$$f(x) = \frac{e^{x} - e^{-x}}{e^{x} + e^{-x}} \tag{7}$$

and found that the number of out-of-gamut incidences have reduced significantly.

3.4 Test Results

Although we have performed our techniques on more than one set of data, the results for only one set are shown here. Images 1 and 2 show two original images, forestl and forest2. Image 3 is the result obtained by applying a one-stage mapping to correct it to forest2 while Image 4 was obtained by correcting forest2 to forestl, performed in the rgb colour space. Although the colours in most regions (sky, mountains, trees) in both images are natural and corrected to match each other quite well, the colours belonging to the yellow flowers have gone out-of gamut because they were not represented in the data sets used for constructing the mapping. The sigmoid function from Equation 6 was used to bring the values in forest2 back in-gamut in Image 5, but as can be seen, they haven't been brought back far enough. Equation 7 gave much worse results. When the xyz colour space was used in Image 6, the results were quite improved.

Incremental regression was then used to correct forestl and forest2 towards each other using 4 steps for each. Images 7 and 8 display the results in the second step, and the final results displayed in Images 9 and 10. It is clearly seen that the results have improved gradually at each step. Images 11 and 12 show the best results when both incremental regression and compensation were applied.

4. Conclusion

Multiple polynomial regression provides a good mapping to correct the colour differences in the overlap regions of two adjacent images in an image sequence. When problems occur in the application of this mapping to the full images due to colours not present in the overlap region, we can incorporate representatives from the non-overlap region into the regression. Sigmoid functions and *incremental regression* were also explored and found to reduce significantly the number of out-of-gamut colours.

References

1. Henry R. Kang, Color scanner calibration. *J. of Imaging Science and Technology,* **36**(2):162-170, March/April 1992.
2. J. Schwartz, Color equalization. *Journal of Imaging Science and Technology,* **36**(4):328-334, July/August 1992.
3. H. Haneishi, K. Miyata, H. Yaguchi, and Y. Miyake, A new method for color correction in hardcopy from crt images. *Journal of Imaging Science and Technology,* **37**(1):30-36, Jan/Feb 1993.
4. Katsuhiro Kanamori and Kiroaki Kotera, Color correction technique for hard copies by 4-neighbours interpolation method. *Journal of Imaging Science and Technology,* **36** (1):73-80, Jan/Feb 1992.
5. A. A. Afifi and S. P. Azen, *Statistical Analysis: A Computer Oriented Approach.* Academic Press New York, 1972.
6. William Mendenhall, Richard L. Scheaffer, and Dennis D. Wackerly, *Mathematical Statistics with Applications.* Duxbury Press Boston, 3rd edition, 1986.

Image 1: forest1, original uncorrected image

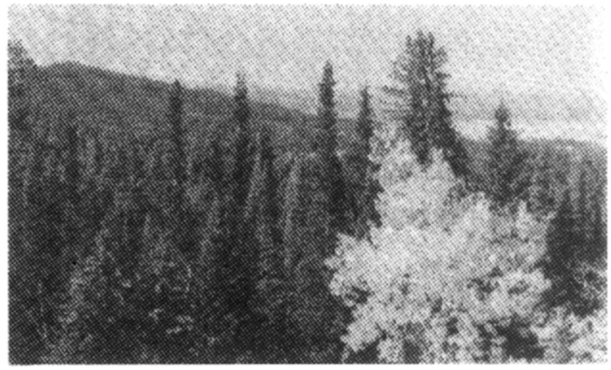

Image 2: forest2, original uncorrected image

Image 3: forest1, corrected to forest2; one step, rgb space

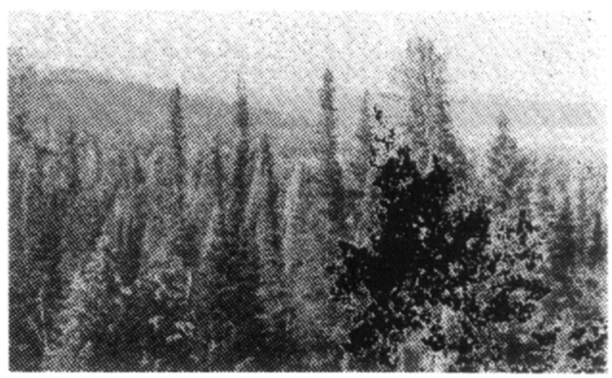

Image 4: forest2, corrected to forest1; one step, rgb space

Image 5: forest2, corrected to forest1; one step, using sigmoid function

Image 6: forest2, corrected to forest1; one step, xyz space

Image 7: Incremental regression; intermediate

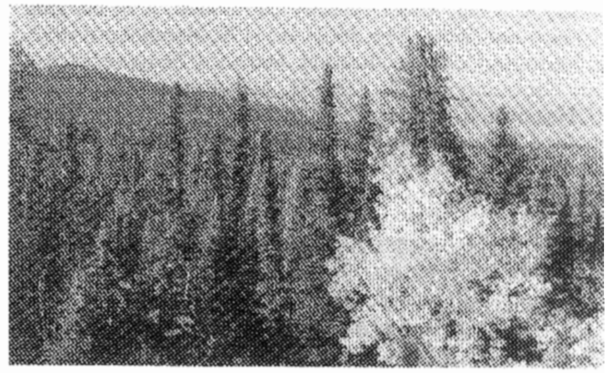

Image 8: Incremental regression; intermediate

Image 9: Incremental regression; final

Image 10: Incremental regression; final

Image 11: Incremental regression with compensation

Image 12: Incremental regression with compensation