Luminance-Based Multi-Scale Retinex

Brian Funt
Kobus Barnard
Michael Brockington
Vlad Cardei
Simon Fraser University

To appear: Proceedings AIC Colour 97 Kyoto 8th Congress of the International Colour Association, May 1997

ABSTRACT

Multi-scale retinex (MSR) processing has been shown to be an effective way to enhance image contrast, but it often has an undesirable desaturating effect on the image colours. A colour-restoration method [2,3]can help mitigate this effect, but our experience is that it simply leads to other problems. In this paper we modify MSR so that it preserves colour fidelity while still enhancing contrast. We then add neural-net based colour constancy processing[7] to this modified version of MSR. The result is a principled approach that provides the contrast-enhancement benefits of MSR and improved colour fidelity.

1. INTRODUCTION

Multi-scale retinex (MSR) processing enhances image contrast [1,2,3] very well. It is particularly good in improving images of scenes where there is a wide range of scene brightnesses, as for example when strong highlights and deep shadows appear in the same image. The major shortcoming of MSR is that it tends to make all image colours greyer than they should be. This greying effect occurs because in MSR the R component (similarly for G and B) of each pixel is replaced with the ratio of its value to the average R component of its neighboring pixels. Wherever the image colour is relatively constant, a pixel's colour will be similar to the average of its neighbors' colours so the ratio in all three channels will be one–in other words it will look grey.

Recent MSR work [2,3] has tried to overcome this greying effect by introducing a colour-restoration step to replace some of the colour lost during the contrast-enhancement step. We find that this step creates more problems than it solves. The exact amount of colour that should be put back depends on how much was removed, and this is a function of the scene and the parameters used. Hence the method tends to make the images more colourful, but not necessarily using the best colours.

MSR evolved from work on human colour constancy [4-6], and it is put forth as a method for obtaining partial colour constancy. However, the colour-constancy method implicit in MSR is not particularly effective since it is primarily a version of the grey-world, colour-balancing algorithm. Furthermore, it gets undermined by the subsequent colour-restoration step.

We propose a simple solution to these difficulties. The main idea is to separate the dynamic-range-compression component of MSR from the colour-constancy component. Furthermore, the dynamic range compression is done so that only the image luminance is changed–the chromaticity or hue of each pixel remains unchanged.

2. SUMMARY OF MULTI-SCALE RETINEX

MSR is explained easily from single-scale Retinex. For SSR we have [2,3].

$$R_i(x,y,c) = \log\{I_i(x,y)\} - \log\{F(x,y,c) * I_i(x,y)\} \tag{1}$$

where $R_i(x,y,c)$ is the output for channel *i*, $I_i(x,y)$ is the image value for channel *i*, $*$ denotes convolution, and $F(x,y,c)$ is the Gaussian function:

$$F(x,y,c) = Ke^{-(x^2+y^2)/c^2} \tag{2}$$

with *K* selected so that:

$$\iint F(x,y,c)\,dx\,dy = 1 \tag{3}$$

In the above, the constant *c* is the scale. The MSR output is simply the weighted sum of several SSR's with different scales:

$$R_{M_i}(x,y,\mathbf{w},\mathbf{c}) = \sum_{n=1}^{N} w_n R_i(x,y,c_n) \tag{4}$$

where $R_{M_i}(x,y)$ is the MSR result for channel *i*, $\mathbf{w} = (w_1, w_2, \ldots, w_N)$ where w_n is the weight of the n^{th} SSR, $\mathbf{c} = (c_1, c_2, \ldots, c_N)$, where c_n is the scale of the n^{th} SSR, and we insist that $\sum_{n=1}^{N} w_n = 1$. In [2] the authors state that the choice of scales is application dependent and that for most applications at least three scales are required. Equal weighting of the scales is usually adequate. The example illustrated in Figure 1 of [2] uses scales of 15, 80, and 250 pixels.

The result of the above processing will have both negative and positive RGB values, and the histogram will typically have large tails. Thus a final gain-offset is applied as described in more detail in [1].

As mentioned above, this processing can cause image colours to go towards grey, and thus an additional processing step is proposed in [3]:

$$R'_{M_i}(x,y,\mathbf{w},\mathbf{c},C) = R_{M_i}(x,y,\mathbf{w},\mathbf{c}) * I'_i(x,y,C) \tag{5}$$

$I'_i(x,y,C)$ is given by:

$$I'_i(x,y,C) = \log\left(1 + C\frac{I_i(x,y)}{\sum_{i=1}^{3} I_i(x,y)}\right) \tag{6}$$

where we have taken the liberty to use log(1+x) in place of log(x) to ensure a positive result.

3. THE MODIFIED ALGORITHM

Each step of the above processing is problematic with respect to maintaining colour fidelity. Each operation in the above sequence changes the image colours. The logarithm in (1) applied to each of the 3 channels independently creates a colour shift. The differencing step in (1) moves the image colours towards grey. Finally, the colour-restoration step multiplies the result by the logarithm of the original colour, which changes the colour in a way which is hard to characterize. More specifically, the restoration effect is a non-linear function of the original image colour and the processed image colour, itself a function of the original image. The amount of colour added with this scheme can at best only approximate the colour removed in the first step; this confounds any colour constancy processing that may have been intended.

Of course, it may be the case that the colour balance of the input image is incorrect, and should be changed. This occurs when there is a mismatch between the illumination for which the imaging system is calibrated and the actual scene illumination. In this case, colour constancy processing is required, and our approach is to apply a sophisticated colour constancy algorithm to the image to estimate the proper image chromaticities. If it is known in advance that the

chromaticities are correct, than this step can be omitted. The colour constancy method we used for this study is the approach described in [7] which corrects for the mismatch between the camera and an unknown illuminant using a neural network.

The second step of the algorithm is to apply MSR processing to the image luminance as follows. We map the input intensity $I_I = \sum_i I_i$ (in the case of three-channels $I_I = I_{red} + I_{green} + I_{blue}$) to the output intensity, $R_I = \sum_i R_i$ using formula (1) which becomes:

$$R_I(x, y, c) = \log\{I_I(x, y)\} - \log\{F(x, y, c) * I_I(x, y)\} \tag{7}$$

with F(x,y,c) given by (2) above. To get a luminance version of MSR, we simply use formula (4) with the arbitrary channel *i* being replaced by the single-intensity result. The next step is to apply the gain-offset method described in [1] to the luminance. Thus having determined the desired luminance, we set each pixel to have the same chromaticity as in the input by:

$$R_i(x, y) = R_I(x, y) \frac{I_i(x, y)}{\sum_i I_i(x, y)} \tag{9}$$

This step can produce pixel values above 255. One possible solution is simply to scale the range to fit. However, often a better result is obtained by applying the gain-offset to the upper range. There will be a slight error in the chromaticities of the pixels that are clipped, but this is not normally noticeable.

4. RESULTS

We have tested our modified MSR method on a number of images; however, we cannot reproduce them in colour in these proceedings. Rather than attempt to portray colour results in black and white, we have made the results of a controlled sequence of images available on the Internet [8]. In that sequence we took images of the same scene with a shadow of varying strengths using two very differently coloured lights. The first was a regular incandescent bulb which is a good illuminant for the 3200°K setting on our CCD camera. The second illuminant was a cool white fluorescent behind a blue filter, which creates an illuminant similar in chromaticity to that of deep blue sky. The 3200°K camera setting was always used, creating a blue cast in the images under the simulated skylight in order to test the colour constancy facility of the revised MSR algorithm.

In one image the incandescent light source was near the camera resulting in an image with good colour balance and devoid of shadows. This was used as a reference. By moving the lights, shadows of increasing strength were cast across the images. In order to explore the method fully, for each illuminant an image with an extraordinarily dark shadow was taken by combining several images taken at different apertures.

In general, the original MSR method greyed out the images. We used the colour-restoration scheme to adjust the colour to match the reference image in order to determine empirically for a reasonable value of C in equation (6). For this image C was 100. Since we are interested in automatic processing we do not wish to go through this procedure for every image. To compare standard MSR to our method which has no such parameter, we feel it is fair to leave the value of C at the empirically discovered value of 100. Since the modified MSR is designed to preserve colour, it did not grey out the image. It, therefore, did not require colour restoration.

We turn now to the images with the blue colour cast. Here the original MSR moves the images towards grey, and somewhat towards the appropriate colour, achieving some degree of colour constancy. The colour is still far from the standard. When the colour restoration was applied, using the same constant as above, the image colours moved back towards the original,

incorrect colour. In fact, it is hard to see how to fix this problem with the original colour restoration method, even if one is allowed to change the parameter manually.

In the case of modified MSR, the colour constancy processing using the method describe in [7] works well, producing an image close to the desired coloured as defined by the standard image. The subsequent MSR luminance processing preserves this colour, producing an image which has both the benefits of the MSR dynamic range compression as well as correct colours.

5. CONCLUSION

The original MSR mixes colour constancy with dynamic range compression, making it difficult to get consistently good colour results. In this work we combine neural-network-based colour constancy with modifications to MSR so that it operates on luminance and preserves colour. This provides both the dynamic range benefits of MSR processing and colour fidelity.

6. ACKNOWLEGEMENTS

This work was supported by the Natural Sciences and Engineering Research Council of Canada and by a gift from Hewlett-Packard Incorporated.

7. REFERENCES

1. Zia-ur Rahman, "Properties of a Center/Surround Retinex Part One: Signal Processing Design," NASA Contractor Report #198195, 1995.
2. Zia-ur Rahman, Daniel J. Jobson, and Glenn A. Woodell, "A Multiscale Retinex for Colour Rendition and Dynamic Range Compression," *SPIE International Symposium on Optical Science, Engineering and Instrumentation, Applications of Digital Image Processing XIX*, Proceedings SPIE 2825, Andrew G. Tescher, ed., 1996.
3. Daniel J. Jobson, Zia-ur Rahman, and Glenn A. Woodell, "Retinex Image Processing: Improved Fidelity To Direct Visual Observation," *Proceedings of the IS&T/SID Fourth Colour Imaging Conference: Colour Science, Systems and Applications, Scottsdale, Arizona*, November, pp. 124-126, 1996.
4. E. H. Land, “Recent advances in Retinex theory and some implications for cortical computations: Colour vision and the natural image,” *Proc. Natl. Acad. Sci.*, 80, pp. 5163-5169, 1983.
5. E. H. Land, “Recent advances in Retinex theory,” Vision Res., 26, pp. 7-21, 1986.
6. Edwin H. Land, “An alternative technique for the computation of the designator in the Retinex theory of colour vision,” *Proc. Natl. Acad. Sci.* USA, Vol. 83, pp. 3078-3080
7. Funt, B., Cardei, V. and Barnard, K., “Learning Colour Constancy,” Proc. Fourth IS&T/SID Colour Imaging Conf., pp. 58-60, Scottsdale, Nov. 19-22, 1996.
8. ftp.cs.sfu.ca/pub/cs/colour/AIC-97

AUTHORS' ADDRESSES

Brian FUNT, Kobus BARNARD, Michael BROCKINGTON, and Vlad CARDEI
School of Computing Science, Simon Fraser University, 8888 University Drive, Burnaby, British Columbia, Canada V5A 1S6, (funt, kobus, iis, cardei)@cs.sfu.ca