

Resiliency of the Multiscale Retinex Image Enhancement Algorithm

Zia-ur Rahman[†], *Daniel J. Jobson*[‡] and *Glenn A. Woodell*[‡]
[†]*College of William & Mary*, [‡]*NASA Langley Research Center*

Abstract

The multiscale retinex with color restoration (MSRCR) continues to prove itself in extensive testing to be a very versatile automatic image enhancement algorithm that simultaneously provides dynamic range compression, color constancy, and color rendition. However, issues remain with regard to the resiliency of the MSRCR to different image sources and arbitrary image manipulations which may have been applied prior to retinex processing. In this paper we define these areas of concern, provide experimental results, and, examine the effects of commonly occurring image manipulations on retinex performance. In virtually all cases the MSRCR is highly resilient to the effects of both the image-source variations and commonly encountered prior image-processing. Significant artifacts are primarily observed for the case of selective color channel clipping in large dark zones in an image. These issues are of concern in the processing of digital image archives and other applications where there is neither control over the image acquisition process, nor knowledge about any processing done on the data beforehand.

Introduction

The Multiscale Retinex¹ (MSR) is a generalization of the single-scale retinex² (SSR), which, in turn, is based upon the last version of Land's center/surround retinex³. The current version, the multiscale retinex with color restoration (MSRCR), combines the dynamic range compression and color constancy of the MSR with a color 'restoration' filter that provides excellent color rendition⁴⁻⁶. The MSRCR has been tested on a very large suite of images. However, concerns about its resiliency to both artifacts owing to digital image formation, and, to the digital processing performed on the image prior to the application of the MSRCR need to be addressed. We provide a general overview of the types of operations that can be performed on the image prior to dissemination and discuss their effect on the MSRCR output.

Resiliency

Webster's Collegiate Dictionary defines *resiliency* as the "ability to to recover from or adjust easily to misfortune or change." We have applied the MSRCR to images where we have no information either about the process that was used to form the image, or about any processing algorithms that were applied to the image. Resiliency in this context refers to the ability of the MSRCR to produce *good* (visual) images regardless of the characteristics of input image. Figure 1 shows the original image* that we use throughout this paper and the MSRCR output using 4 scales. Though there appears to be a "graying-out" of the bright areas when compared with the original image, the sharpness and visibility of detail in the MSRCR output, more than compensate for any lack of local contrast. We use this original image and pre-process it to to simulate the commonly applied "enhancement" filters. Results are shown later in the paper.

Multiscale Retinex with Color Restoration

The general form of the MSRCR can be summarized by the following equation:

$$\mathcal{R}_{M_i}(x, y) = G_r F_i(x, y) \sum_{s=1}^S w_s (\log [I_i(x, y)] - \log [I_i(x, y) * M_s(x, y)]) - O_r, \quad i = 1, \dots, N \quad (1)$$

where \mathcal{R}_{M_i} is the i th band of the MSRCR output, S is the number of scales being used, w_s is the weight of the scale, I_i is the i th band of the input image, and N is the number of bands in the input image. The surround function M_s is defined by

$$M_s(x, y) = K \exp \left[-\frac{\sigma_s^2}{x^2 + y^2} \right],$$

where σ_s is the standard deviation of the s th surround function, and $\iint K \exp \left[-\frac{\sigma_s^2}{x^2 + y^2} \right] dx dy = 1$; $F_i(x, y)$

*Courtesy of the NASA Johnson Space Center.



Original



MSRCR with N=4

Figure 1: The source image for all the simulations and the MSRCR output



Original with negative offset



MSRCR output



"Corrected" input



MSRCR of corrected input

Figure 2: MSRCR resiliency to the presence of negative offsets.

are the color restoration functions defined by

$$F_i(x, y) = G_f \log \left[\frac{I_i(x, y)}{\sum_{n=1}^N I_n(x, y)} - O_f \right].$$

G_r and O_r are, respectively, the final gain and offset values needed to scale the output of the log domain operations to the (R,G,B) color space, and G_f and O_f control the degree to which the color restoration function $F(x, y)$ affects the overall color of the output image. These constants, the number of scales, S , and the widths of the surround functions, σ_s , are image independent[†] in the sense that we apply the same (canonical) set of constants to every image that we process.

Image formation and image processing related issues

Digital images can either be directly acquired with digital cameras, or can be obtained through scanners from prints, negatives and slides. All of these devices have built-in automatic functions for conversion from the analog to the digital domain, to provide modest dynamic range compression, and to correct for the film transfer characteristics in the case of scanners, and for filtering certain wavelengths in the case of cameras. In addition there are typically manual color balance controls. The exact implementation of these functions is generally device dependent, but their overall effect is directly observable in the output image. Resiliency is of significant interest here because for most images obtained from, say, the Internet, we neither know the the image was acquired, nor do we know the type of pre-processing it has undergone. What this means that we do not have access to the *scene* from which the digital image was acquired, and we have to be able to deduce the source of artifacts and correct for them because they affect the overall visual quality of the retinexed image.

Commonly occurring operations performed on images are:

Negative Offset

The most common effect that we have encountered is the presence of a strong negative offset in the image. The minimum value below a threshold is pegged to black[‡]. This is an attempt to increase the dynamic range (i.e. visual contrast) provided by the device but is often photometrically incorrect and results in *false zeroes*. The effect on the MSRCR is to produce a harsher-than-normal contrast. A more extreme case of this, also often encountered, is signal clipping where low signal information is actually lost.

[†]Typically for 512×512 images. The σ_s may change with the dimensions of images.

[‡](0,0,0) in the (Red,Green,Blue) coordinates.

When this effect is severe, the MSRCR produces much stronger color saturations, since the overall effect of the negative offset is to increase the relative strength of signals between color channels. Particularly strong effects are observed when setting individual band values below a certain threshold to zero leaves one or two color bands with a non-zero value, thus fundamentally changing the color at that location. Since the MSRCR produces a log spatial/spectral ratio, this situation, in effect, represents a "divide-by-zero" condition that can lead to significant color artifacts. For instance, if this happens in large dark zones in the image, it often manifests itself as *neon* streaking of shocking color.

Figure 2 shows the original image from Figure 1 with a negative offset applied to it. As can be seen by comparing the two figures, the contrast is better in Figure 2, but the effect on the MSRCR output is also severe. Though it is very evident in the gray-scale images shown in this paper, the MSRCR output in this case has become overly harsh.[§] A simple correction, i.e. application of a positive offset to the original image can mitigate this effect and is shown in the bottom row of Figure 2.

Automatic Gain and Offset

Auto gain can be performed either in hardware at device level, or in software as part of the drivers/application packages that read the images from the hardware. In auto gain/offset operations, a negative offset is typically applied to map the minimum value to black and then a gain is applied to map the resultant maximum value to white. Care must be taken to ensure that actual white exists in the scene. The MSRCR is very resilient to such adjustments. Since the difference between the MSRCR outputs in the original and the auto/gain case is insignificant, the result is not shown here.

Positive Offset:

Typically brightness in an image is increased by applying a positive offset, i.e. the mean value of the image is increased. This often manifests itself as an overall haziness in the input image. Though the application of the MSRCR reduces this haziness, there is still a sense of haziness overall. Further alleviation of this effect can be achieved by reducing the final offset value O_r (Equation 1) from its canonical value. An alternate way to improve the output is by applying a negative offset to the original image before the application of the MSRCR. It should be noted that an overall haziness in the output of the MSRCR is a good indication of the presence of positive offsets in the original image. The MSRCR output for either of these methods is essentially the same.

[§]For color images you may access a copy of the paper from <ftp://vipsun.larc.nasa.gov/retinex/retpubs/>.



Original with positive offset



MSRCR output



"Corrected" input



MSRCR of corrected input

Figure 3: Resiliency of the MSRCR to positive offsets in the original image.



Original with gamma=2.5



MSRCR output

Figure 4: Application of gamma correction increases the overall dynamic range that can be displayed but has an overall hazy effect on the image.

Figure 3 shows the original image with a positive offset. Again, the MSRCR provides more detail in the dark regions than the input image, though the contrast is not as good as that shown in Figure 1. The second row of Figure 3 shows the effect of applying the MSRCR to the corrected image. Most of the dynamic range shown in the original MSRCR is preserved though at a slight loss of contrast. The color images make this point more clearly.

Non-linear gamma correction:

The dynamic range of the image is adjusted using non-linear gamma correction to compensate for the too-dark and too-bright regions. Mathematically,

$$O_i(x, y) = [I_i(x, y)]^{\frac{1}{\gamma}},$$

where O , and I are the output and input respectively. The MSRCR is quite resilient to this non-linearity over a range of $0.5 \leq \gamma \leq 1.8$, though it is more resilient to changes for $\gamma \leq 1.0$. The primary effect of applying $\gamma > 1.0$ is similar to that obtained when positive offsets are present, i.e. overall hazy appearance (Figure 4). The haziness from the application of gamma correction can be reduced in a similar manner to that used for images containing positive offsets.

Lossy compression:

Lossy compression is often applied to images both to allow more images to be archived, and for faster distribution over the Internet. Depending upon the type of the algorithm, the effects of lossy compression can manifest themselves as block-edge artifacts, overall loss in resolution, i.e. crispness of edges, or a loss in dynamic range. The extent to which these artifacts are ‘enhanced’ is extremely dependent on the image content—the MSRCR is a context-based algorithm—but is most marked in large dark zones. Generally though the retinex produces more visual information along with the JPEG artifacts, so an image-specific trade-off occurs where the benefits must be weighed against the quality required for a specific application.

Figure 5 shows the effects of applying the MSRCR to a JPEG’d image. Again, though not very clear in the gray-scale images show here, observe the block artifacts that are enhanced in the top right corner of the MSRCR output. Also note the increased dynamic range that is evident. We have noted that whereas the application of the MSRCR to lossy compressed images tends to enhance the artifacts introduced by the compression algorithm, the application of the compression algorithm to the MSRCR output does not suffer from similar problems. The bottom row of Figure 5 shows an image where the compression takes place after the application of the MSRCR. Care has been taken so

that the MSRCR file and the original JPEG file are almost identical in size. It is evident though that the application of the compression algorithm *after* the application of the MSRCR does not suffer from the same artifacts as those shown in the top row of Figure 5.

There are other issues that arise when dealing with heavily compressed images, but that is a topic for another paper!

Conclusions

We have provided a brief description of the commonly encountered ‘problems’ introduced inevitably in a digital image due to the nature of the acquisition process and the pre-processing algorithms. Since in many image enhancement applications—e.g. images obtained from the Internet—we neither know the source of the image (digital camera or scanner), nor do we know how the images have been ‘enhanced,’ it is critical that we understand the effects of these common processes on the output of the MSRCR. We recognize that in such cases, slight modifications to the canonical set of constants may need to be made in order to obtain the best possible visual quality. However, though the presence of these operations in the input image can adversely affect the overall visual quality of the output image produced by the MSRCR, even the ‘not-the-best’ MSRCR output is still typically better than the original image in terms of contrast, visual quality, and color constancy. The MSRCR has thus proven to be quite resilient to many of the arbitrary operations that are used in digital image formation and can thus be truly considered a fully automatic process.

References

- [1] Z. Rahman, D. Jobson, and G. A. Woodell, ‘Multiscale retinex for color image enhancement,’ in *Proceedings of the IEEE International Conference on Image Processing*, IEEE, 1996.
- [2] D. J. Jobson, Z. Rahman, and G. A. Woodell, ‘Properties and performance of a center/surround retinex,’ *IEEE Trans. on Image Processing: Special Issue on Color Processing*, vol. 6, pp. 451–462, March 1997.
- [3] E. Land, ‘An alternative technique for the computation of the designator in the retinex theory of color vision,’ *Proc. Nat. Acad. Sci.*, vol. 83, pp. 3078–3080, 1986.
- [4] D. J. Jobson, Z. Rahman, and G. A. Woodell, ‘A multi-scale Retinex for bridging the gap between color images and the human observation of scenes,’ *IEEE Transactions on Image Processing: Special Issue on Color Processing*, vol. 6, pp. 965–976, July 1997.
- [5] Z. Rahman, D. Jobson, and G. A. Woodell, ‘Multiscale retinex for color rendition and dynamic range compression,’ in *Applications of Digital Image Processing XIX* (A. G. Tescher, ed.), Proc. SPIE 2847, 1996.
- [6] D. Jobson, Z. Rahman, and G. A. Woodell, ‘Retinex image processing: Improved fidelity for direct visual observation,’ in *Proceedings of the IS&T Fourth color Imaging Conference: Color Science, Systems, and Applications*, pp. 124–126, IS&T, 1996.



JPEG image



MSRCR output of JPEG image



MSRCR image



JPEG output of MSRCR image

Figure 5: MSRCR tends to enhance JPEG artifacts but the application of the MSRCR before compression can lead to better results.