Analysis and Improvement of Multi-Scale Retinex

Kobus Barnard and Brian Funt

Abstract

The main thrust of this paper is to modify the

enhancement so that the processing is more justified from a theoretical standpoint. This leads to a new

flexible, maintains color fidelity, and still preserves the contrast-enhancement benefits of the original MSR

implicit processing goals of MSR. By decoupling the MSR operations from one another, we build an algorithm composed of independent steps that separates out the issues

method. We then extend MSR with color constancy and chromaticity-preserving contrast enhancement.

Recent work [1,2,3,4] advocates MSR as a method ofimage enhancement which provides color constancy and

dynamic range compression. Nonetheless, there are aThe chief conceptual problem is that a number ofimage-process

sensitive to color.

processing goals, depending on the circumstances:

1) Compensating for uncalibrated devices (gamma correction)
2) Color constancy processing
3) Local dynamic range compression
4) Global dynamic range compression
5) Color enhancement

intertwined, and as a result, the colors are changed in image dependent and unpredictable ways. We will disentangle these tasks, and develop a sound theoretical basis for them. In addition, when the

proceed relative to a well defined color baseline, and such a baseline is provided by our approach. To do this we use an effective, neural-network-based, color constancy algorithm [5] to correct for mismatches between the illuminant and the camera balance. Then we use a modified version of MSR which does not change

baseline for subsequent color enhancement if required by the application.

Overview of the original MSR method

MSR is explained easily from single-scale Retinex [2, 3, 4]. For SSR we have:

(1)

where $R_i(x,y,c)$ is the output for channel "i",

is the image value for channel "i", ⊗ denotes

is a Gaussian surround function explicitly given by:

(2)

with K selected so that:

(3)

In the above, the constant "c" is the scale. The MSR

with different scales:

(4)

where is the MSR result for channel "i", where W

n'th SSR, is the scale

of the n'th SSR, and we insist that

the authors state that the choice of scales is application dependent, but that for most applications at least three scales are required, and that equal weighting is usually adequate. The example illustrated in Figure 1 of [2] uses

used in [4].

negative and positive RGB values, and the histogram will typically have large tails. Thus a final gain-offset

detail below.

This processing can cause image colors to go towards gray, and thus an additional processing step isproposed in [1]:

(5)

where $I_i'(x,y,C)$ is given by:

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

where we have taken the liberty to use log(1+x) in place

125 is suggested for C; for [6] we empirically settled on a value of 100 for a specific test image. The difference between using these two values is small. In formula (5) of [4] a second constant is used which is simply a multiplier of the result: .
However, in our implementation this constant is absorbed in the final gain-offset step.

A few more words about the final gain-offset step

needed to have good contrast, as the resultant image

of consistently removing these tails led to the pair of gain-offset constants recently published in [4][1]. However, our experience prior to that publication did not leave us confident that the best choice is particularly image independent, especially for images from various

clipping point based on clipping a few percent of the pixels on either side. The extrema of the clipping points from each of the three channels is used, in order to apply the same gain-offset to all three. We have also explored clipping a fixed percentage of the range, as well as looking for image independent gain offset parameters, as suggested in [4]. All methods work to some extent, but we do not have an appropriate metric for performance, and additional work is required in this area.

Unfortunately, for the purposes of comparison, the methods and constants used for the gain-offset can substantially affect image appearance. This is especially

interest has been compressed with a logarithm operation. In addition, the resulting color is sensitive to the
Formula (6) of [4] seems to imply that none is needed, but in our implementation it is necessary to get a reasonable result. Furthermore, the color is sensitive to the associated constants.

MSR and Color fidelity

and output devices. In particular, we require a linear relationship between scene radiance and CRT luminance of the three channels (up to a uniformly applied

CRT's gamma must be taken into account. The standard power-law method of gamma correction is only

from monitor to monitor. As a result, we have found it worthwhile to calibrate our monitor with a spectra-radiometer.

There is an interesting relationship with the original formulation of MSR and gamma correction. That method uses a channel-independent logarithm, which normally would have the side effect of changing

color shift is far less of a problem when displaying the result without gamma correction than would be expected. In fact, applying gamma correction to the result of MSR processing normally gives poor results. Specifically the images look washed out and over gamma corrected. The problem with just accepting and using this coincidence as a conveniently provided

correction) is meant to compensate for devices, but now one is committed to a single method, and thus the result is device dependent. Regardless, since MSR can to some extent play the role of gamma correction, it is

MSR results on a monitor.

Device calibration is also an issue on input. It is interesting to note that given a histogram or range based

MSR method can be quite resilient to insufficient information about whether a gamma correction has been applied to the input. The gamma becomes roughly a constant factor due to the channel independent logarithm

based gain-offset adjustment. For confirmation, we have verified that there is little visual difference between MSR output of gamma corrected and non gammacorrected

The second color problem with the original MSR approach is that the image colors tend to be desaturated grayish. This is due to the manner in which gray-world-based color constancy processing is applied to relatively small image neighbourhoods. Each pixel's color is compared to the average of the colors in a surrounding neighbourhood.
means that the MSR result will tend towards gray regardless of the color of the region.

---

[1] We have not yet been able to obtain reasonable results using the formula and constants in [4].

Later versions of MSR include a processing step which puts back some of the color that was removed.The intermediate image colors are modified by a non-linear function of the original image colors. This processing has the obvious problem of changing the predict. A second problem with the color correction step a camera balanced for a redder light, will be too blue. MSR without color correction will move the color of the wall towards gray, and thus achieve some degree ofcolor constancy. However, if the color correction step is now used, the color of the wall will be moved back towards blue!

Another color problem with standard MSR processing is complement color bleeding at certain color edges due to the local contrast enhancement. Consider a white card mounted on a yellow background. For simplicity, consider that the red and the green channels of the yellow are similar to that of white, and the blue is substantially smaller. Then only the blue channelwill change at the boundary, and the blue channel of the white near the boundary will be enhanced relative to the others which represent neutral. Hence the white card will have a blue halo near the boundary.

Color preserving MSR

We now outline an alternative approach to MSR. As processing goals/effects of MSR so that each one can be linear. Then we optionally apply color constancy processing followed by MSR style processing to an appropriately defined image luminance. The processing here can take many forms, of which two are discussed in detail below. The RGB of the output image pixels arethen set so that their chromaticity is the same as in the (possibly color corrected) linear input image, but their luminances are the result of the previous processing increasing the color saturation, can be applied. Finally, the image is mapped into the appropriate space to give linear output on the target device. In the case of a CRT monitor, this can be approximated by a gamma correction. We now provide some additional details.

As discussed above, color fidelity is best achieved ifthe input is proportional to scene radiance. Thus we attempt to linearize the input if this is not the case. We have experimented with input from a Sony CCD camera as well as Kodak photo CD images. In the case of the camera, we have verified that it is linear. We linearize [7]. It is not known how well this corresponds to the radiance in the original scenes, but for the purposes of experimentation we assume it is linear.

If color constancy is an issue for the application, it is dealt with next. For the purposes of this paper, wedefine co which the imaging system is calibrated and the actual illuminant chromaticity of the scene. Color correction illuminant independent description of the scene. Most assume that the input is linear, and thus a good result is dependent on the linearity considerations discussed above. In fact, using the above definition for color linear space.

Standard MSR has its roots in the latest color constancy work by Land [8,9,10], and color constancy processing is one of the purported goals of MSR processing. However, the color constancy processing inherent in standard MSR processing has several weaknesses. First, it attempts to do color correction in a non-linear space. Second, it essentially is based on the except that there are better algorithms available (see, for example, [5,11,12,13]). A more serious problem is that the implementation of the gray world algorithm is not optimal. Color constancy algorithms generally make chromaticity varies spatially (the most common assumption being that it is uniform), and then exploit that assumption. In the case of MSR, the use of a large scale implies some confidence that the illumination uniformity is wide, but the use of smaller scales yieldspoor colo image. Averaging the results mitigates the errors, but also reduces the chances for good performance, and thus is unsatisfactory. We posit that if illumination uniformity is an issue, it should be dealt with explicitly in the algorithm (as is done in [14]). Otherwise, the as this gives the most effective color constancy processing.

The color constancy algorithm used for our chromaticity of the scene illuminant [5]. This is then illuminant for the imaging system. The performance of this algorithm is significantly better than gray world based methods.

The next step is to apply MSR style processing on an appropriately defined expression of the image method is simpler and changes the image less, and may be preferable for images from sources known to have small dynamic range. The second method is designed to approximate the dynamic range compression of the original MSR method. The method is that it is more appropriate on images with high dynamic range. In order to investigate therelationship of the various methods and input dynamic

range we created some images with extended dynamic range by either combining a number of images taken at different apertures, or averaging a large number of images.

For the first method we apply MSR style processing *without*

luminance defined by

channels $I_{I=}\ I_{red} + I_{green} + I_{blue}\ )$ as follows. For each scale we map the input intensity to the outputintensity,

, using formula (1) where without

logarithms the subtraction becomes a division:

$$R_I(x,y,\ ) = I_I(x,y)/F(x,y,c) \otimes I_I(x,y) \quad (7)$$

with $F($ 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. This method has the appeal that the luminance is in a space which is locally approximately linear, and thus the image which require little or no change should look more natural.

With an appropriate choice of scales, the above compression. This is the case because a very small scale

will remove all intensity differences, and reduce the

the above method to images with large dynamic range often gives a poor result at sharp shadow edges. The

but the edge itself becomes a dark area between two light areas, and thus looks unnatural.

The reason for this difference is that a large part of the

the logarithm operation. This can be verified by applying the processing without any ratios. The observation that the logarithm operation has a definite benefit leads us to the second method for luminance based MSR style processing.

This method is designed to provide the same dynamic range compression as original MSR. Here wedefine the image luminance by the geometric mean of

$I_{I=} \left(\prod_i^N I_i\right)^{1/N}$. Although it is possible to use the arithmetic mean (as was done in [6]), the geometric mean is intuitively superior, as it gives a cleaner correspondence between the luminance ofstandard MSR and the luminance based alternative.

Having computed the luminance, standard MSR processing is now applied to it, this time including the logarithm operation. In order to obtain an output luminance comparable to standard MSR, an additional step is needed. This is due to the observation above that MSR output should not be gamma corrected. Since we wish to gamma correct the output of the modifiedalgorithm

MSR luminance result. Again the correspondence between the effect and the desired result is better served by the use of the geometric mean in place of the

only dealing with luminance, the reverse gamma correction need not be exact, and is adequately implemented with a power function. Specifically we raise the luminance to the 2.2 power, with any power in monitor. If even more dynamic range compression is

required, it can be obtained by simply omitting the reverse gamma correction step, but images processed in this manner to tend to look unnatural.

The next step is to apply the histogram based (or other) gain-offset method described above to the luminance. Thus having determined the desired relative intensify, we set each channel to the same chromaticity

$$R_i(x,y) = R_I(x,y)\frac{I_i\ x,y)}{I_I\ x,y)} \quad (8)$$

The processing so far has been designed to maintain color fidelity. However, this is not the same as producing the most pleasing color. If color enhancement is desired, then it is best added at this stage. For example, for some applications, increasing color saturation may be desired.

Next we map the pixels into the output range, typically [0, 255], recalling that the zero point is already set by the bottom clipping of the intensity. One possible solution is to simply scale the range to fit. However, often a better result is obtained by allowing some clipping of the upper range. The chromaticities of the pixels that are clipped will be a slightly incorrect,but this

recommended, however, to do the same with the bottom

amount of clipping on the bottom by doing so when the luminance range is adjusted.

The final step of the algorithm is to map the output into a space which produces linear output on the target device. In the case of a CRT monitor, this may be approximated by gamma correction.

In summary, we have an algorithm which

standard MSR, but is precise with respect to color. In

we only need to perform convolutions on the luminance. Even if convolutions are performed using Fourier transforms, this is a non-negligible saving.

Resu

We have tested the modified method of MSR processing on a number of images. Rather than attempt to portray the results available We first verified that for standard images, the first

form of the dynamic range compression usually gives reasonable results. These images included ones from Kodak photo CD and ones taken with a three chip, 8

images had sufficiently strong shadow boundaries that