

# Local Color Correction Using Non-Linear Masking

*Nathan Moroney  
Hewlett-Packard Laboratories  
Palo Alto, California*

## Abstract

Tone reproduction is a key component in image quality and historically tone correction has been performed on a global basis for a given device or image. Several recent investigations have used a variety of algorithms to demonstrate that local color correction provides a significant improvement in image quality relative to a global correction. This paper presents a local color correction operation that is based on non-linear masking. The operation is very fast and does not require any manual adjustments, such as those involved in dodging and burning. The algorithm is roughly equivalent to deriving a specific tone reproduction curve for each pixel in the image. An example image is processed using the proposed masking scheme and the resulting effect is discussed. Finally, the relationship to CIECAM97s is explored and it is proposed that dynamically setting the background parameter results in similar local color corrections.

## Introduction

Tone reproduction is a critical aspect of the perceived quality of a reproduction.<sup>1</sup> Historically tone corrections have been global one-dimensional operations in which one input value maps to one and only one output value.<sup>2</sup> This approach results in a reasonable correction if the original dynamic range is fairly close to the dynamic range of the reproduction. For a much larger dynamic range original it becomes increasingly difficult to design a global tone correction that will accommodate both shadow and highlight detail.

Therefore, there has been ongoing research<sup>3-15</sup> in the area of local color correction and processing high dynamic range images. A local color correction will provide a method to map one input value to many different output values, depending on the values of neighboring pixels. This allows for simultaneous shadow and highlight adjustments. It is beyond the scope of this paper to provide a comparison of the previously proposed algorithms for local color correction. However, some limitations that have been encountered in some cases are computational complexity, lack of calibration procedure, "halo" artifacts, iterative processing, segmentation failures, lack of direct control

parameters, incompatibility with look up table processing or requiring manual adjustments.

This paper demonstrates a local color correction technique that is simple, fast and comparable to the other local color correction methods.<sup>3-15</sup> This method uses non-linear masking in order to perform local color correction.<sup>16</sup> It should be noted that initial investigations of this method of local color correction were not based on any model of visual perception. To illustrate, traditional unsharp masking provides a powerful tool to sharpen images but care should be taken not to attribute the functionality of the algorithm to specific aspects of the human visual system. The proposed masking algorithm provides a simple, computationally efficient and easily parameterized framework for performing local color corrections.

Local color correction is also useful for enhancing or improving certain types of images. To illustrate, conventional photography may yield an image in which one region of the image is underexposed while another region is correctly exposed. Flash photography is another instance in which one region in the image may be over-exposed while another is under-exposed. Correcting for failures during image capture or for illumination fall-off are not, strictly speaking, the goals of a vision model. Aesthetic and contextual issues further complicate the issue of local color correction but will not be considered in this paper.

## Basic Algorithm

The first part of the proposed algorithm consists of the derivation of an image mask. The second part of the algorithm is the combination of the mask and the input image. This is shown in flowchart form in Figure 1 where one operator computes the mask and the other combines the mask and the input image. In this paper, the mask is simply an inverted low-pass filtered monochrome version of the input image. The combination operation is then a variable power function where the exponent is computed from the mask.

The mask can be computed by taking an image and converting it to a monochrome image. This monochrome image is then inverted and finally the image is blurred using a large radius filter. The mask is monochrome in order to avoid distorting the chroma of the image. The image is inverted to indicate which regions of the image

will be lightened or darkened. For example, a dark region in the image will have a light mask value and will therefore be lightened. However, a positive mask can also be used as will be discussed in a following section. The image should be blurred such that image features can no longer be recognized. For instance, a 72 dpi monitor image could be blurred using a gaussian blur with a radius of 15 pixels. If the mask is not blurred, then the image contrast will be excessively reduced. In comparison, if the mask is overly blurred then this algorithm reduces to simple gamma correction. A computationally efficient implementation can use an image thumbnail or a significantly sub-sampled image in order to compute the mask.

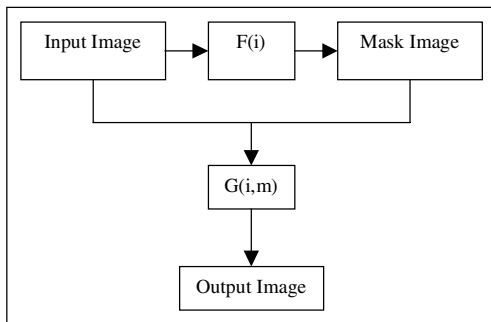


Figure 1. Basic flowchart for non-linear masking.

The combination operation consists of a simple power function where the exponent is computed using the mask value. Mask values greater than 128 will result in an exponent less than 1 while mask values less than 128 will yield exponents greater than 1. Mask values of 128 will result in a mask value of 1 and will not change the input data. Example curves showing how input values are mapped to output values as a function of the mask value are shown in Figure 2. This operation essentially consists of performing a pixel-by-pixel gamma correction. This can be written as the following equation:

$$Output = 255 * \left( \frac{Input}{255} \right)^{2 \left( \frac{128 - Mask}{128} \right)} \quad (1)$$

Equation 1 can be implemented as a fully populated two dimensional table where the input and mask values are the input and the output is the pre-calculated value for that combination. Furthermore, the equation can easily be modified to adjust the identity value or the range of the exponents. This would be useful for a limited parameter adjustment based on user preferences or even image properties. This tone reproduction surface can be thought of as a two-dimensional version of a tone reproduction curve.

There are a number of intriguing variations on the mask derivation and combination operators. However, the method described demonstrates one of the simplest techniques for performing local color correction using non-

linear masking. Notice that while the mask derivation process may be similar to that used during the traditional process of unsharp masking, the non-linear combination operation is much different than the additive operation used during unsharp masking and therefore results in local color adjustments.

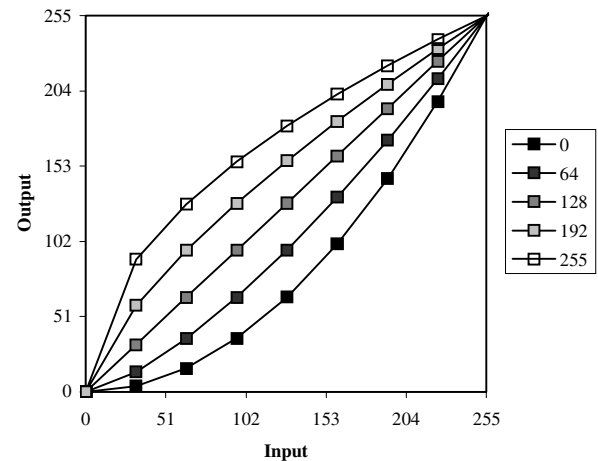


Figure 2. Example curves showing how input 8 bit RGB values are mapped to output values as a function of the mask value.

## Results

The results of local color correction using masking are shown in Figures 3. This figure shows an original black and white image on the top. The figure in the image is slightly underexposed but a global tone curve to correct the figure would reduce the highlight detail in the sky. Below the original image is the mask as computed using the previously described technique. It is interesting to note the similarity of the mask image to a negative afterimage.<sup>17</sup> The bottom is the corrected image. Note that light regions in the image correspond to dark regions in the mask and are darkened in the corrected image. Conversely, dark regions in the image correspond to light regions in the mask and are lightened in the corrected image.

The fence behind the boy in Figure 3 is changed without any overshoot or ringing around the edges. In addition, the image has a white border and even this strong edge does not introduce any distortions. This is a result of two factors. First, the non-linear masking is bounded at the extremes so that the white and black points of the image are never adjusted regardless of the mask values. This means that the correction will be less at the extremes, where haloing may be more common. Second, the low pass nature of the mask means that tone corrections will occur slowly over the image. These gradual transitions result in local tone adjustments that are not obvious for complex images or regions. For simple images or regions on the other hand the mask is simple and in the extreme reduces to a simple global tone correction.



Figure 3. Original, mask and locally color corrected image shown from top to bottom. The original image has a white border surrounding the image.

### Discussion

Equation 1 is very similar to the CIECAM97s<sup>18</sup> formula for lightness. In this case J or lightness is a function of the luminance of the background, as well as other parameters. Fairchild notes that “specification of the background is absolutely necessary for modeling simultaneous contrast”.<sup>19</sup> A simplified form of CIECAM97s lightness can be written:

$$J = 100 \left( \frac{A}{A_w} \right)^{c_1} \left( 1 + c_2 \left( \frac{Y_b}{Y_w} \right)^{1/2} \right) \tag{2}$$

where c1 and c2 are viewing condition parameters that result in an exponent that ranges from 0.69 to about 1.38 for an average surround. The form of this equation with the normalization, exponent and scaling is very similar to equation 1.

Equation 2 yields a family of lightness curves for a range of background luminances. This is shown in figure 4 and can be compared with figure 2. The CIECAM97s local color correction uses a positive mask and therefore the order of white to black backgrounds is reversed. Equation 1 can also be modified to assume a positive mask but the author finds working with a negative mask is useful for assessing where the image will be lightened and darkened by looking at the mask. The positive mask is used for CIECAM97s because of the perceptual significance of Yb.

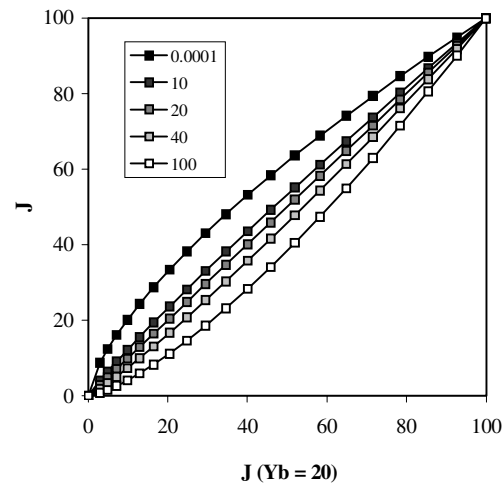


Figure 4. Family of lightness curves for various CIECAM97s background or Yb settings. The lightness for a given background is plotted against the lightness for Yb of 20.

A non-linear masking operation similar to that previously described can be performed using CIECAM97s. The images are first converted to CIECAM97s J, a<sub>97</sub>, b<sub>97</sub> coordinates and then a positive lightness mask is computed using the J channel of the image. The resulting images are similar to those previously generated using just an exponential operator on the RGB channels. This demonstrates that local color correction can be computed using

CIECAM97s and the result is an image in which highlights can be darkened and shadows lightened in one step. This processing does not require segmentation or iteration and is based on the results of psychophysical observations. Furthermore, using CIECAM97s instead of a device RGB provides a method of calibrating the local color correction.

There are numerous possible applications of this type of processing. One application is to better map high dynamic range scenes to lower dynamic range display and output devices. Another application is to correct for non-uniform illumination, such as flash photography, backlit subjects or flat-bed scans of three-dimensional objects. This type of processing can also be used given an original that has regions that are incorrectly exposed. This type of color correction can also be used for aesthetic purposes, such as to reduce differences in illumination when compositing

## Conclusion

Local color correction can be performed using non-linear masking. This correction can simultaneously lighten shadows and darken highlights and is based on a simple pixel-wise "gamma" correction of input data. Using a low pass inverted monochrome version of the original will balance global and local contrast changes and reduce chroma distortions. This algorithm can be used for images with uneven exposure or flash illumination. It can also be used for object scanning and adjusting high dynamic range scene data during scene rendering. While there are a number of alternative algorithms for performing local color correction, it is interesting to note the relative simplicity of the proposed algorithm and its close relationship to the lightness scale that is incorporated in the CIECAM97s color appearance model. This implies that it may be useful to consider constructing a color appearance model from near to far.

Local color correction would seem to be a complementary technology to traditional pixel based processing. For example, once the image has been corrected it can then be reproduced on a given printer using a traditional pixel based pipeline. Use of a neighborhood image processing technique for one part of the imaging pipeline does not exclude subsequent pixel based image processing. This suggests that pixel and neighborhood imaging operations could be used where appropriate to optimize the overall quality and speed of the imaging pipeline.

## References

1. R.W.G. Hunt, *The Reproduction of Color*, Fountain Press, England, (1987).
2. T. H. James, Ed., *The Theory of the Photographic Process*, Eastman Kodak Company, Rochester NY, 536-561 (1977).
3. Jack Tumblin and Greg Turk, "LCIS: A Boundary Hierarchy for Detail Preserving Contrast Reduction", *SIGGRAPH 99*, 83-90 (1999).
4. Sumanta N. Pattanaik, James A. Ferwerda, Mark D. Fairchild and Donald P. Greenberg, "A multiscale model of

adaptation and spatial visual for realistic image display", *SIGGRAPH 98*, 287-298 (1998).

5. G. W. Larson, H. Rushmeir, and C. Piatko, "A visibility matching tone reproduction operator for high dynamic range scenes", *IEEE Trans. Vis. And Comp. Graph.* **3**, 291-306 (1997).
6. Paul E. Debevec and Detendra Malik, "Recovering High Dynamic Range Radiance Data from Photographs", *SIGGRAPH '97*, 369-378 (1997).
7. Paloma G. Diaz, MiniGuia de Aprendizaje rapido Adobe PhotoShop, INFORBOOK'S, S.L., 196-197, (1995).
8. Shigeo Sakaue, et al "Adaptive gamma processing of the video cameras for the expansion of the dynamic range", *IEEE Transactions on Consumer Electronics*, **41**, 555-562 (Aug 1995).
9. Tae Keun Kim, et al, "Contrast enhancement system using spatially adaptive histogram equalization with temporal filtering", *IEEE Transactions on Consumer Electronics*, **44**, 82-87 (Feb 1998).
10. Yu Wang, et al, "Image enhancement based on equal area dualistic sub-image histogram equalization method", *IEEE Transactions on Consumer Electronics*, **45**, 68-75 (1999).
11. Jeffrey M. DiCarlo and Brian A. Wandell, "Rendering high dynamic range images", *Proc. SPIE, Image Sensors*, **3965**, (2000).
12. Toshikazu Kato and Yuichi Kobayashi, "A High Fidelity Contrast Improving Method Based on Human Vision Mechanisms", *Proc. IEEE Int Conf Multimedia Comp and Systems Vol II* (1998).
13. Jia-Guu Leu, "Image Contrast Enhancement Based on the Intensities of Edge Pixels", *Computer Graphical Models and Image Processing*, **54**, 497-506, (1992).
14. Edwin H. Land and John McCann, "Lightness and the Retinex Theory", *J. Opt. Soc. Am.*, **61**, 1-11, (1971).
15. Zia-Ur Rahman, Glen A. Woodell, and Daniel J. Jobson, "A Comparison of Multi-Scale Retinex with Other Image Processing Techniques", *Proc. IS&T 50<sup>th</sup> Anniversary Conference*, (1997).
16. The author has filed a patent application with the U.S. Patent Office on local color correction using non-linear masking.
17. Cristina A. Burbeck, "Negative afterimages and photopic luminance adaptation in human vision", *J. Opt. Soc. Am. A*, **3**, 1159-1165 (1986).
18. M.R. Luo and R.W.G. Hunt, "The Structure of the CIE 1997 Colour Appearance Model (CIECAM97s)", *Color Res. App.* **23** 138-146 (1998).
19. Mark D. Fairchild, *Color Appearance Models*, Addison-Wesley, Reading Massachusetts, (1998).

## Biography

Nathan Moroney is a color scientist in the Color Imaging and Printing Technologies Department of Hewlett-Packard Laboratories. Previously, he worked for the Barcelona division of Hewlett-Packard and at the RIT Research Corporation. He has a Bachelors degree in color science from the Philadelphia University and a Masters degree in color science from the Munsell Color Science Laboratory of RIT. His research interests include color appearance models and color variability. He is currently the technical chair for CIE TC8-01.