

# Contrast Enhancement of Images using Human Contrast Sensitivity

Aditi Majumder, Sandy Irani \*

Computer Science Department, University of California, Irvine

## Abstract

Study of contrast sensitivity of the human eye shows that our contrast discrimination sensitivity follows the weber law for suprathreshold levels. In this paper, we apply this fact effectively to design a contrast enhancement method for images that improves the local image contrast by controlling the local image gradient. Unlike previous methods, we achieve this without segmenting the image either in the spatial (multi-scale) or frequency (multi-resolution) domain.

We pose contrast enhancement as an optimization problem that maximizes the average local contrast of an image. The optimization formulation includes a perceptual constraint derived directly from human suprathreshold contrast sensitivity function. Then, we propose a greedy heuristic, controlled by a single parameter, to approximate this optimization problem. The results generated by our method is superior to existing techniques showing none of the common artifacts of contrast enhancements like halos, hue shift, and intensity burn-outs.

**CR Categories:** I.4.0 [Image Processing and Computer Vision]: General—[image displays]; I.4.8 [Image Processing and Computer Vision]: Scene Analysis—[color photometry]; H.1.2 [Models and Principles]: User/Machine Systems—[human factors]

**Keywords:** Human Perception, Contrast Sensitivity, Contrast Enhancement

## 1 Introduction

The sensitivity of the human eye to spatially varying contrast is a well-studied problem in the perception literature and has been studied at two levels: threshold and suprathreshold. Threshold contrast sensitivity studies the minimum contrast required for human detection of a pattern, while suprathreshold contrast studies the perceived contrast when it is above the minimum threshold level. These studies show that contrast discrimination sensitivity can be quantified with a single parameter, especially at suprathreshold levels [Valois and Valois 1990; Whittle 1986]. In this paper, we use the suprathreshold contrast sensitivity to design a new contrast enhancement technique for 2D images.

The problem of enhancing contrast of images enjoys much attention and spans a wide gamut of applications, ranging from improving visual quality of photographs acquired with poor illumination [Oakley and Satherley 1998] to medical imaging [Boccignone and Picariello 1997]. Common techniques for *global* contrast enhancements like global stretching and histogram equalization do not always produce good results, especially for images with large spatial variation in contrast. To address this issue, a large number of *local* contrast enhancement methods have been proposed that use some form of image segmentation either in the spatial (multi-scale) or frequency (multi-resolution) domain followed by the application of different contrast enhancement operators on the segments. These approaches differ in the way they generate the multi-scale or multi-resolution image representation, or in the contrast enhancement operators they use to enhance contrast after segmentation. Image segmentation has been achieved using methods such as anisotropic diffusion [Boccignone and Picariello 1997], non-linear pyramidal techniques [Toet 1992], multi-scale morphological techniques [Toet 1990; Mukhopadhyay and Chanda 2002], multi-resolution splines [Burt and Adelson 1983], mountain clustering [Hanmandlu et al. 2001] or retinex theory [Munteanu and Rosa 2001; Rahman et al. 1996]. Contrast enhancement of the segments has been achieved using morphological operators [Mukhopadhyay and Chanda 2002], wavelet transformations [Velde 1999], curvelet transformations [Stark et al. 2003], k-sigma clipping [Munteanu and Rosa 2001; Rahman et al. 1996], fuzzy logic [Hanmandlu et al. 2000; Hanmandlu et al. 2001] and genetic algorithms [Shyu and Leou 1998].

In this paper we present a local contrast enhancement method driven by an objective function that is controlled by a single parameter derived from the suprathreshold contrast discrimination sensitivity of the human eye. The perception of contrast is directly related to the local luminance difference i.e. the local luminance gradient at any point in the image. Our goal is to enhance these gradients. Methods dealing with gradient manipulation need to integrate the gradient field for image reconstruction. This is an approximately invertible problem, achieved by solving the Poisson equation, and has been used recently to achieve contrast enhancement and seamless image editing [Fattal et al. 2002; Prez et al. 2003]. However, these methods are often cumbersome to implement because they involve differential equations dealing with millions of variables. Instead, we achieve gradient enhancement by treating images as height-fields and processing them in a way that can be controlled by the single parameter. We pose this as an optimization problem that maximizes the local average contrast in an image strictly guided by a perceptual constraint derived directly from the human suprathreshold contrast discrimination sensitivity. In addition, the range of the color values are strictly constrained to avoid artifacts due to saturation of colors. To solve this optimization problem we propose a new greedy iterative algorithm. We compare the results from this algorithm with existing different global and local contrast enhancement techniques and show that our results are superior than any traditional or state-of-the art contrast enhancement techniques. By imposing explicit constraints in

\*e-mail: {majumder,irani}@ics.uci.edu

our optimization formulation, we are able to avoid all common artifacts of contrast enhancement like halos, intensity burn-out, hue shift and introduction of noise.

## 2 Suprathreshold Contrast Sensitivity

In this section, we derive the equation that guides the sensitivity of the human eye to brightness differences at different intensities. Contrast detection has been studied in vision perception literature for decades [Valois and Valois 1990]. Threshold contrast sensitivity functions (CSF) define the minimum contrast required to detect a sinusoidal grating of a particular mean and spatial frequency. These are bow shaped plots with peak sensitivity at about 5-6 cycles/degree and the frequency for peak sensitivity decreases as mean brightness decreases.

So far we have talked about threshold CSF. But most of our everyday vision is at suprathreshold (above threshold) levels. Recently there has been a large number of work to study the contrast discrimination sensitivity of human beings for suprathreshold vision. Of this, we are particularly interested in the study of contrast increments in the context of our contrast enhancement application. [Whittle 1986] presents one of the most comprehensive studies in this direction. This shows that for suprathreshold contrast  $C$ , contrast discrimination threshold follows the Weber law, i.e.

$$\frac{\partial C}{C} = \lambda \quad (1)$$

where  $\tau$  is a constant. This indicates that for visible contrast enhancement, higher contrast patterns need higher contrast increments. This forms the mainstay of our contrast enhancement method.

But, before we use the above equation, we need to generalize it for different spatial frequencies. A recent study [Kingdom and Whittle 1996] showed that the character of contrast discrimination is similar for both sinusoidal and square waves of different spatial frequencies. This finding is corroborated by other works [Barten 1999; Georgeson and Sullivan 1975] confirming that the suprathreshold contrast discrimination characteristics show little variation across spatial frequencies. Also, [Peli 1990; Wilson 1991] has shown the contrast perception to be a quasi-local phenomenon, mainly because we use our foveal vision to estimate local contrast.

Using the above, we derive a simple equation for contrast enhancement of images. We define the local contrast of the image to be proportional to the local gradient of the image. In other words,

$$C \propto \frac{\partial I}{\partial x} \quad (2)$$

where  $I(x, y)$  is the image,  $C$  is the contrast and  $\lambda$  is the constant of proportionality. Equation 1 indicates that to achieve the same perceived increase in contrast across an image, larger gradients have to be stretched more than smaller gradients. In fact, the stretching should be performed in such a fashion that the contrast increment is proportional to the initial gradient. Thus,

$$\frac{\partial I'}{\partial x} \geq (1 + \lambda) \frac{\partial I}{\partial x} \quad (3)$$

where  $I'(x, y)$  is the contrast enhanced image. Using the above facts, we express the contrast enhancement of an image  $I(x, y)$  by a single parameter  $\tau$  as

$$1 \leq \frac{\frac{\partial I'}{\partial x}}{\frac{\partial I}{\partial x}} \leq (1 + \tau) \quad (4)$$

where  $\tau \geq \lambda$ . The lower bound assures that contrast *reduction* does not occur at any point in the image and the upper bound assures that the contrast enhancement is bounded. [Mantiuk et al. 2006] have shown the constant  $\lambda$  to be close to 1 by fitting a curve to the experimental data of [Whittle 1986]. Thus contrast enhancement in the images will only be visible for  $(1 + \tau) \geq 2$  assuring that the Equation 3 is satisfied. Equation 4, though simple, is very effective in practice to achieve contrast enhancement of images.

## 3 The Method for Gray Images

We pose the local contrast enhancement problem as an optimization problem. We design a scalar optimization function derived from Equation 2 that captures the overall contrast of an image, and seek to maximize it subject to the constraint described by Equation 4. In addition, we also constrain the color range of the output image to avoid over or under saturation artifacts.

### 3.1 Optimization Problem

First, we formulate the contrast enhancement optimization problem for gray images. We consider the intensity values of a grayscale image to be representative of the luminance values at the pixel locations.

We pose the optimization problem as follows. We propose to maximize the objective function

$$f(\Omega) = \frac{1}{4|\Omega|} \sum_{p \in \Omega} \sum_{q \in N_4(p)} \frac{I'(p) - I'(q)}{I(p) - I(q)} \quad (5)$$

subject to a *perceptual constraint*

$$1 \leq \frac{I'(p) - I'(q)}{I(p) - I(q)} \leq (1 + \tau) \quad (6)$$

and a *saturation constraint*

$$L \leq I'(p) \leq U \quad (7)$$

where scalar functions  $I(p)$  and  $I'(p)$  represent the gray values at pixel  $p$  of the input and output images respectively,  $\Omega$  denotes sets of pixels that makes up the image,  $|\Omega|$  denotes the cardinality of  $\Omega$ ,  $N_4(p)$  denotes the set of four neighbors of  $p$ ,  $L$  and  $U$  are the lower and upper bounds of the gray values (e.g.  $L = 0$  and  $U = 255$  for images that have gray values between 0 and 255), and  $\tau > 0$  is the single parameter that controls the amount of enhancement achieved. This objective function is derived from Equation 2 as a sum of the perceived local contrast over the whole image, expressed in the discrete domain. Note that it also acts as a metric to quantify the amount of enhancement achieved. The perceptual constraint (Equation 6) is derived directly from Equation 4 by expressing it in the discrete domain. The lower

bound in this constraint assures two properties: the gradients are never shrunk; the sign of the gradients are preserved. Finally, the saturation constraint (Equation 7) ensures that the output image does not have saturated intensity values. Note that the saturation constraint does not control the gradient but just the range of values a pixel is allowed to have. Thus the pixels in the very dark or very bright regions of the image will still have their gradients enhanced.

### 3.2 Greedy Iterative Algorithm

We propose an iterative, greedy algorithm to try to maximize the objective function above subject to the constraints. Being local in nature, our method adapts to the changing local contrast across the image achieving different degrees of enhancement at different spatial locations of the image.

Our algorithm is based on the fundamental observation that given two neighboring pixels with gray values  $r$  and  $s$ ,  $r \neq s$ , scaling them both by a factor of  $(1 + \tau)$  results in  $r'$  and  $s'$  such that

$$\frac{r' - s'}{r - s} = (1 + \tau) \quad (8)$$

Thus if we simply scale the values  $I(p)$ ,  $\forall p \in \Omega$ , by a factor of  $(1 + \tau)$ , we obtain the maximum possible value for  $f(\Omega)$ . However, this could cause violation of Equation 7 at some pixel  $p$ , leading to saturation of intensity at that point. To avoid this, we adopt an iterative strategy, employing a greedy approach at each iteration.

We consider the image  $I$  as a height-field (along the Z axis) sampled at the grid points of a  $m \times n$  uniform grid (on the XY plane). This set of samples represents  $\Omega$  for a  $m \times n$  rectangular image. Thus, every pixel  $p \in \Omega$  is a grid point and the height at  $p$ ,  $I(p)$ , is within  $L$  and  $U$ .

For each iteration, we consider a plane perpendicular to the Z axis at  $b$ ,  $L \leq b \leq U$ . Next, we generate a  $m \times n$  matrix  $R$  by simple thresholding of  $I$  and identifying the regions of the height field  $I$  which are above the plane  $b$  as

$$R(i, j) = \begin{cases} 1 & \text{if } I(i, j) > b \\ 0 & \text{if } I(i, j) \leq b \end{cases} \quad (9)$$

$R$  generates a graph where the vertices are those pixels  $(i, j)$  such that  $R(i, j) = 1$  and two vertices are adjacent if they are neighbors in the image. We then identify connected components in this graph. Each such component, represented by  $h_i^b$ , is called a *hillock*; the subscript denotes the component number or label and the superscript denotes the plane used to define the hillocks. Next, the pixels in each hillock are scaled up by an amount such that no pixel belonging to the hillock is pushed beyond  $U$  or has the gradient around it enhanced by a factor of more than  $(1 + \tau)$ . The scaling factor is chosen individually for each hillock.

Our method involves successively sweeping threshold planes  $b_i$  such that  $L \leq b_i < U$  and at each sweep, greedily scaling the hillocks respecting the constraints. Note that as we sweep successive planes, a hillock  $h_i^b$  can split into  $h_j^{b+1}$  and  $h_k^{b+1}$  or remain unchanged, but two hillocks  $h_s^b$  and  $h_t^b$  can never merge to form  $h_u^{b+1}$ . This results from the fact that our threshold plane strictly increases from one sweep to the next and hence the pixels examined at a stage are a subset of the pixels examined at previous stages. Thus, we obtain

the new hillocks by only searching amongst hillocks from the immediately preceding sweep.

For low values of  $b$ , the size of the hillocks are large. Hence, the enhancement achieved on hillocks might not be close to  $(1 + \tau)$  because of the increased chances of a peak close to  $U$  in each hillock. As  $b$  increases, the large connected components are divided so that smaller hillocks can be enhanced more than before.

This step of sweeping planes from  $L$  to  $U$  pronounces only the local hillocks of  $I$  and the image thus generated is denoted by  $I_1$ . However, further enhancement can be achieved by enhancing the local *valleys* also. Thus the second stage of the our method applies the same technique to the complement of  $I_1$  given by  $U - I_1(p)$ . The image generated from the second stage is denoted by  $I_2$ .  $I_2$  is then complemented again to generate the enhanced output image  $I' = U - I_2(p)$ .

### 3.3 Performance Improvement

We perform  $U - L$  sweeps to generate each of  $I_1$  and  $I_2$ . In each sweep we identify connected components in a  $m \times n$  matrix. Thus, the time-complexity of our algorithm is theoretically  $O((U - L)mn)$ . However, we perform some optimizations to reduce both the space and time complexity of the method.

We observe that hillocks split at local points of minima or saddle point. So, we sweep planes only at specific  $b_i$ s where some points in the height field attain a local minima or saddle point. This helps us to achieve an improved running time complexity of  $O(sm)$  where  $s$  is the number of planes swept (number of local maximas, local minimas and saddle points in the input image). This idea is illustrated in Fig 1. However, note that this example is constructed to illustrate the method and we have exaggerated the enhancements for better comprehension. In practice, many images have numerous local minima and saddle points. The result is that the threshold usually only increases by one or two values in 8-bit greyscale images. This results in a process that is more time-intensive than necessary. Therefore, we have an additional parameter  $\Delta$  which is a lower bound on the amount by which the threshold must increase in consecutive iterations. This, in effect, skips some of the sweeping planes. For greyscale images whose values are in the range from 0 to 255, a  $\Delta$  of 5 or 10 still produces excellent results. This reduces the value of  $s$  in the running time to be at most  $255/\Delta$ . The results are compared in Figure 2.

We also observe that disjoint hillocks do not interact with each other. So, to make our method memory efficient, we process each hillock in a depth first manner before proceeding to the next hillock.

To summarize, following is the pseudocode of our algorithm.

Algorithm *Enhance*( $\tau, I, L, U$ )

Input: Control parameters  $\tau$  and  $\Delta$   
Input Image  $I$   
Lower and upper bounds  $L$  and  $U$   
Output: Enhanced Image  $I'$

**Begin**

1.  $I' \leftarrow I$ ;
2.  $I' = \text{ProcessHillocks}(I', \tau, \Delta)$ ;
3.  $I' \leftarrow U - I'$ ;

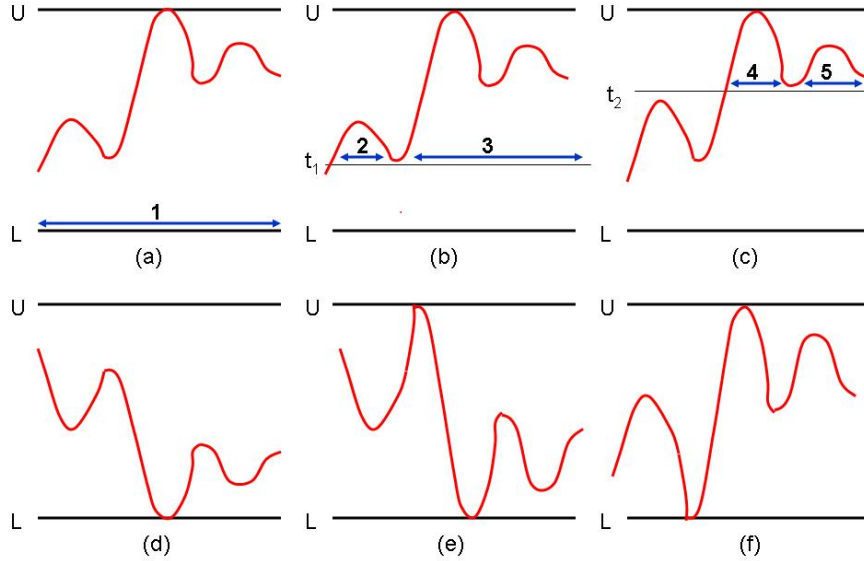


Figure 1: Graphs showing some key steps in our algorithm when applied to a 1D signal. In (a), the threshold is zero and there is a single hillock comprised of all the pixels in the input. The hillock is stretched so that the maximum pixel value reaches saturation. In Figure (b), the threshold is increased to  $t_1$  and hillock 1 is now divided into Hillocks 2 and 3. Hillock 3 can not be stretched any further since its highest pixel is already saturated. However, Hillock 2 can be stretched so that the local enhancement of each pixel (as denoted in Equation 6), reaches  $1 + \tau$ . Since this is the maximum enhancement that can be achieved, no further processing is performed on Hillock 2. In Figure (c), the threshold is  $t_2$  and Hillock 3 splits into Hillock 4 and 5. Only Hillock 5 can be further enhanced since Hillock 4 has a saturated pixel. Hillock 4 is stretched so that the local enhancement of each pixel reaches  $1 + \tau$ . In the second pass, the image from Figure (c) is inverted to produce Figure (d). Hillocks are processed and stretched as in the first pass to produce Figure (e). Image (e) is then inverted back to obtain the final enhanced image in (f).

4.  $I' = \text{ProcessHillocks}(I', \tau, \Delta)$ ;  
5.  $I' \leftarrow U - I'$ ;  
6. **Return**  $I'$ ;  
**End**

Algorithm *ProcessHillocks*( $I, \tau, \Delta$ )  
Input: Input Image  $I$   
Control parameters  $\tau$  and  $\Delta$   
Output: Image  $I'$

**Begin**

1.  $I' = I$ ;
2. Create empty Stack  $S$  of Hillocks
3. Create hillock  $h$
4. Initialize pixels in  $h$  to be all pixels in  $I$
5. Initialize threshold( $h$ ) to be 0
6. Push  $h$  onto  $S$
7. While  $S$  not empty repeat
8.  $h = S.\text{pop}()$
9. Find connected components of pixels in  $h$  whose value is at least threshold( $h$ )
10. For each connected component  $c$
11. Create new hillock  $h'$
12. Initialize pixels in  $h'$  to be all pixels in  $c$
13. For each pixel  $p$  in  $h'$
14.  $I'(p) = (1 + s) * (I'(p) - t) + t$   
where  $t$  is threshold( $h$ ) and  $s$  is the maximum value over the entire hillock such that none of the constraints are violated.
15. Let threshold( $h'$ ) be the minimum of  $I(p)$  over all pixels  $p$  in  $h'$  that are local minima or

saddle points and  $I(p)$  is at least threshold ( $h'$ )  
18. threshold( $h'$ ) =  $\max\{\text{threshold}(h'), \text{threshold}(h) + \Delta\}$   
17. Push  $h'$  onto  $S$ .  
**End**

*Enhance* calls the main routine *ProcessHillocks* on the original image and then on the inverted image so that hillocks get pushed upwards and valleys get pushed downwards. *ProcessHillocks* maintains a stack of hillocks. Each hillock maintains a set of pixels which is disjoint from the pixels in any other hillock. Each hillock also maintains a threshold parameter. In each iteration, the top hillock is popped and the threshold is applied to all the pixels in the hillock. These pixels whose value is above the threshold generate an underlying graph with edges between neighboring pixels. We then compute the connected components of this graph and create a new hillock for each component. In Step 14, all the pixels in each component are then stretched upwards as much as possible without violating any of the predefined constraints. Thresholds are moved upwards and all the resulting hillocks are pushed onto the stack.

### 3.4 Results

Fig 2 shows the result of applying our method to low-contrast gray images for different values of  $\tau$ . We studied the effect of skipping some of the sweep planes by increasing  $\Delta$  and found that we can increase the performance by at least an order of magnitude before seeing visible differences.




Figure 2: The original gray image (a), enhanced gray image using  $\tau$  of 0.3 (b) and 2 (c). Note that parts of the image that have achieved saturation for  $\tau = 0.3$  do not undergo anymore enhancement or show any saturation artifact for higher  $\tau$  of 2. Yet, note further enhancement of areas like the steel band on the oxygen cylinder, the driver's thigh and the pipes in the background. (c) is generated by sweeping all planes, (d) and (e) are generated by sweeping one of every five and fifty of  $(U - L)$  planes respectively. Note that despite having five times fewer number of sweep planes, (d) is perceptibly indistinguishable from (c). (e) can be distinguished from (c) by the lower contrast background and upper leg of the diver. (f) is the same image enhanced using fattal's method of stretching gradients using a Poisson solver. Compare this with (c) which is enhanced using our method with  $\tau = 2$ . (f) shows noise artifacts (better visible when displayed in true resolution). In addition, our method in (c) shows much better performance in driving the darker regions to black achieving better contrast in the image.



Figure 3: The original gray image (a), enhanced using global histogram equalization (b), enhanced using local histogram equalization (c) and enhanced using our method with  $\tau = 1$ . Note that global histogram equalization leads to oversaturation of parts of the image in (b). While local histogram equalization alleviates that problem, it ends up introducing noise in the background and changes the appearance of parts of the image in (c) like the shirt. Our method in (d) does not suffer from both of these and achieves an image which is closer to the original in its appearance.




Figure 4: The result of our method on some medical images. An x-ray image (a) enhanced with  $\tau = 1$  (b), and the image of some stem cells taken through a microscope (c) enhanced using  $\tau = 2$  (d).

Figure 2 illustrates this.

Figure 3 compares our method with standard techniques for contrast enhancement that uses global and local histogram equalization respectively. Figure 2 compares our method on gray images with the recent method proposed in [Fattal et al. 2002] that stretches the gradient image directly and then generates the enhanced image from the modified gradient field using a poisson solver. Figure 4 shows the result of our method on some medical images.

With the ideal parameter of  $\Delta = 1$ , our optimized code takes about 10 seconds to process a  $500 \times 500$  image. However, by setting  $\Delta = 10$ , we can process the same image in about a couple of seconds.

### 3.5 Evaluation

The advantage of our formulation of the contrast enhancement problem as an optimization problem lies in the fact that the objective function, defined in Equation 5, can be directly used as a metric to evaluate the amount of *average contrast enhancement* (ACE) achieved across the whole image. Note that according to the constraints of the optimization problem, the maximum average contrast that can be achieved *without respecting saturation constraints* is given by  $1 + \tau$ . The saturation constraints restrict the actual enhancement achieved to less than or equal to  $1 + \tau$  thus making the image free of any artifacts. However, as  $\tau$  increases, the effect of the saturation constraint becomes stronger since larger number of pixels reach saturation when being enhanced and hence needs to be restricted not to enhance to their fullest extent. Hence, with the increase in  $\tau$ , the ACE achieved falls more and more away from  $1 + \tau$ . Table 1 illustrates the ratio of ACE and  $1 + \tau$  for different images for different values of  $\tau$ . Note that as  $(1 + \tau)$  increases, though the ACE value increases in an absolute sense, the ratio of ACE and  $(1 + \tau)$  decreases as expected.

Table 1 also shows that the same metric can be used to evaluate the effect of different optimizations (like skipping some of the sweeping planes and the reverse pass) on the ACE achieved. Note that increasing the number of skipped planes indeed decreases the ACE. This metric also provides us with interesting insights in the algorithm. For example, it shows that most of the enhancement is achieved in the first pass itself. The enhancement in the second pass is mostly of details, and skipping it can help us improve the performance.

Image	$1 + \tau$	$\Delta$	No. of Passes	ACE	$\frac{ACE}{1 + \tau}$
Diver	2	10	2	1.43	0.715
	3	10	2	1.80	0.6
	4	10	2	2.18	0.545
	4	1	2	2.39	0.598
	4	50	2	1.32	0.33
	4	10	1	2.13	0.532
Blonde	3	10	2	1.45	0.483
	5	10	2	1.84	0.368

Table 1: This table compares the ACE achieved by changing different parameters on which the performance of our greedy algorithm depends.  $(1 + \tau)$  is the parameter controlling the contrast enhancement. Step size defines the gap between the planes which are swept to intersect the luminance field. The number of passes is 2 in both the forward and inverse passes are applied and 1 when only the forward pass is applied.

## 4 Extending to Color Images

The most obvious way to extend the algorithm presented in the preceding section to color images is to apply the method independently to three different color channels, as illustrated in Figure 5. However, this does not assure hue preservation resulting in hue shift, especially with higher values of  $\tau$ . This happens when one of the channels saturates at some pixels of the image while other channels have significant room for enhancement (Figure 6).

To avoid this problem, we take the standard approach of separating the luminance and the chrominance channels of the image (by linearly transforming RGB values to CIE XYZ values) and then applying our method only on the luminance channel.

The XYZ values of the maximum intensity of the primaries, defined by three vectors in the XYZ space,  $\vec{R} = (X_R, Y_R, Z_R)$ ,  $\vec{G} = (X_G, Y_G, Z_G)$  and  $\vec{B} = (X_B, Y_B, Z_B)$ , define the color gamut of the RGB device. The transformation from RGB to XYZ space is defined by a  $3 \times 3$  matrix whose columns correspond to  $\vec{R}$ ,  $\vec{G}$  and  $\vec{B}$ . From the XYZ values at each pixel, we obtain its luminance,  $Y$ , and chromaticity coordinates (defining hue and saturation)  $(x, y)$  [Giorgianni and Madden 1998] by

$$x = \frac{X}{X + Y + Z}; \quad y = \frac{Y}{X + Y + Z}.$$

Next, we perform the enhancement only on  $Y$  keeping the chromaticity coordinates,  $x$  and  $y$ , unchanged. Finally, we convert the  $Yxy$  image back to RGB space via XYZ space. To assure a linear transformation between the RGB and the

XYZ space we apply the standard gamma correction to our images.

However, note that changing the luminance  $Y$  can change the chrominance thereby taking the color out of the gamut of the device leading to saturation artifacts. To avoid this, we modify our saturation constraint that assures that the new color achieved after enhancing  $Y$  is within the color gamut of the RGB device. A color in the XYZ space that can be expressed as a convex combination of  $\bar{R}$ ,  $\bar{G}$  and  $\bar{B}$  is within the device's color gamut. Thus, the saturation constraint can no longer be expressed by a single linear inequality as in Equation 7. Instead, as we enhance  $Y$ , we have to assure that the enhanced color lies within the parallelepiped defined by the convex combination of  $\bar{R}$ ,  $\bar{G}$  and  $\bar{B}$ .

Thus, the luminance enhancement problem can be formulated in a similar fashion as gray scale images. The color at pixel  $p$  given by  $C(p) = (X, Y, Z)$  is to be enhanced to  $C' = (X', Y', Z')$ . The goal is to enhance the luminance  $Y$  to  $Y'$  such that the objective function

$$f(\Omega) = \frac{1}{4|\Omega|} \sum_{p \in \Omega} \sum_{q \in N_4(p)} \frac{Y'(p) - Y'(q)}{Y(p) - Y(q)}, \quad (10)$$

is maximized subject to a *perceptual constraint*

$$1 \leq \frac{Y'(p) - Y'(q)}{Y(p) - Y(q)} \leq (1 + \tau), \quad (11)$$

and a *saturation constraint*

$$(X', Y', Z') = c_R \bar{R} + c_G \bar{G} + c_B \bar{B}, \quad 0.0 \leq c_R, c_G, c_B \leq 1.0, \quad (12)$$

Note that Equation 11 explicitly assures that the chromaticity coordinates are not changed and hence hue is preserved. This optimization is solved applying the same technique as in Section 3 on the luminance channel and by changing the check for saturation constraint as per Equation 12 resulting in hue-preserving contrast enhancement devoid of any hue-shift artifacts (Figure 6).

However, note that for applying the second pass, when we invert the image with respect to a spatially varying saturation envelope, the relative magnitudes of neighboring pixels may not be maintained. This can lead to switching of direction of gradient at some pixels which is undesirable. From our study using the evaluation metric in Section 3.5 we found that second pass leads to insignificant change in ACE. So, for color images we only apply the first pass.

**Results:** The most recent contrast enhancement technique is the one developed by Fattal et. al [Fattal et al. 2002] that does a direct gradient stretching and applies a poisson solver to get the image back from its gradient field. We compare results from our method with this work in Figure 5 and 6. Note the different kinds of artifacts like halo, noise and in particular, hue-shift, which we have avoided entirely. Figure 6 and 7 compare the result of our method with some existing global and local contrast enhancement techniques highlighting the fact that our method is devoid of artifacts.

## 5 Conclusion

In conclusion, we use suprathreshold human contrast sensitivity functions to achieve contrast enhancements of images.

We apply a greedy algorithm to the image in its native resolution without requiring any expensive image segmentation operation. We pose the contrast enhancement as an optimization problem that maximizes an objective function that defines the local average contrast enhancement (ACE) in an image subject to constraints that control the contrast enhancement by a single parameter. We extend this method to color images where hue is preserved while enhancing only the luminance contrast.

The ACE defined by the objective function can act as a metric to compare the contrast enhancement achieved for different methods and different parameters thereof. In addition, the parameter  $\tau$  can be varied spatially over the image to achieve spatially selective enhancement. We have omitted a detailed description of these for lack of space.

Future work in this direction will include exploring the possibility of extending this to video by adding the additional temporal dimension. Since our method treats the image as a height field, it could have interesting applications in terrain or mesh-editing. Finally, implementation of this method on GPUs would allow this method to be used interactively.

## References

- BARTEN, P. G. 1999. Contrast sensitivity of the human eye and its effects on image quality. *SPIE - The International Society for Optical Engineering*, P.O. Box 10 Bellingham Washington 98227-0010. ISBN 0-8194-3496-5.
- BOCCIGNONE, G., AND PICARIELLO, A. 1997. Multiscale contrast enhancement of medical images. *Proceedings of ICASSP*.
- BURT, P. J., AND ADELSON, E. H. 1983. A multiresolution spline with application to image mosaics. *ACM Transactions on Graphics* 2, 4, 217–236.
- FATTAL, R., LISCHINSKI, D., AND WERMAN, M. 2002. Gradient domain high dynamic range compression. *ACM Transactions on Graphics, Proceedings of ACM Siggraph* 21, 3, 249–256.
- GEORGESON, M., AND SULLIVAN, G. 1975. Contrast constancy: Deblurring in human vision by spatial frequency channels. *Journal of Physiology* 252, 627–656.
- GIORGIANNI, E. J., AND MADDEN, T. E. 1998. *Digital Color Management : Encoding Solutions*. Addison Wesley.
- HANMANDLU, M., JHA, D., AND SHARMA, R. 2000. Color image enhancement by fuzzy intensification. *Proceedings of International Conference on Pattern Recognition*.
- HANMANDLU, M., JHA, D., AND SHARMA, R. 2001. Localized contrast enhancement of color images using clustering. *Proceedings of IEEE International Conference on Information Technology: Coding and Computing (ITCC)*.
- KINGDOM, F. A. A., AND WHITTLE, P. 1996. Contrast discrimination at high contrasts reveal the influence of local light adaptation on contrast processing. *Vision Research* 36, 6, 817–829.
- MANTIUK, R., MYSZKOWSKI, K., AND SEIDEL, H.-P. S. 2006. A perceptual framework for contrast processing of high dynamic range images. *ACM Transactions on Applied Perception* 3, 3.

- MUKHOPADHYAY, S., AND CHANDA, B. 2002. Hue preserving color image enhancement using multi-scale morphology. *Indian Conference on Computer Vision, Graphics and Image Processing*.
- MUNTEANU, C., AND ROSA, A. 2001. Color image enhancement using evolutionary principles and the retinex theory of color constancy. *Proceedings 2001 IEEE Signal Processing Society Workshop on Neural Networks for Signal Processing XI*, 393–402.
- OAKLEY, J. P., AND SATHERLEY, B. L. 1998. Improving image quality in poor visibility conditions using a physical model for contrast degradation. *IEEE Transactions on Image Processing* 7, 167–179.
- PELI, E. 1990. Contrast in complex images. *Journal of Optical Society of America A* 7, 10, 2032–2040.
- PREZ, P., GANGNET, M., AND BLAKE, A. 2003. Poisson image editing. *ACM Transactions on Graphics, Proceedings of ACM Siggraph 22*, 3, 313–318.
- RAHMAN, Z., JOBSON, D. J., , AND WOODDELL, G. A. 1996. Multi-scale retinex for color image enhancement. *IEEE International Conference on Image Processing*.
- SHYU, M., AND LEOU, J. 1998. A genetic algorithm approach to color image enhancement. *Pattern Recognition* 31, 7, 871–880.
- STARK, J.-L., MURTAGH, F., CANDÉS, E. J., AND DONOHO, D. L. 2003. Gray and color image contrast enhancement by curvelet transform. *IEEE Transactions on Image Processing* 12, 6.
- TOET, A. 1990. A hierarchical morphological image decomposition. *Pattern Recognition Letters* 11, 4, 267–274.
- TOET, A. 1992. Multi-scale color image enhancement. *Pattern Recognition Letters* 13, 3, 167–174.
- VALOIS, R. L. D., AND VALOIS, K. K. D. 1990. *Spatial Vision*. Oxford University Press.
- VELDE, K. V. 1999. Multi-scale color image enhancement. *Proceedings on International Conference on Image Processing 3*, 584–587.
- WHITTLE, P. 1986. Increments and decrements: Luminance discrimination. *Vision Research* 26, 10, 1677–1691.
- WILSON, H. 1991. Psychophysical models of spatial vision and hyperacuity. *Vision and Visual Dysfunction: Spatial Vision*, D. Regan, Editor, Pan Macmillan, 64–86.






Figure 5: Our method applied to the red, green and blue channel of a color image - the original image (a), enhanced image using  $\tau$  of 1 (b) and 9 (c) - note the differences in the yellow regions of the white flower and the venation on the leaves are further enhanced. The same image enhanced by applying Fattal's method is shown in (d). When comparing with the results of our method in (b) and (c), note the haloing artifacts around the white flower and the distinct shift in hue, especially in the red flower, that changes its appearance. It almost appears that (d) is achieved by hue enhancement of (a), rather than a contrast enhancement of (a).



Figure 6: (a) The original image, (b) our method is applied with  $\tau = 2$  on each of red, green and blue channels independently (note the severe hue shift towards purple in the stairs, arch and wall, and towards green on the ledge above the stairs), (c) the image is first converted to brightness and chrominance channels and our method is applied only to the brightness channel. Note that the hue is now preserved. Compare (c) with results from curvelet transformation [Stark et al. 2003] (d), method based on manipulation of gradient field inverted back using poisson solver [Fattal et al 2002] (e), and method based on multi-scale retinex theory [Rahman et al. 1996] (f). Note that (c) and (d) lead to a noisy image while (d) and (e) change the hue of the image significantly.




Figure 7: (a) The original image, (b) our method with  $\tau = 2$ , (c) multi-scale morphology method [Mukhopadhyay and Chanda 2002] - note the saturation artifacts that gives the image an unrealistic look, (d) Toet's method of multi-scale non-linear pyramid recombination [Toet 1992] - note the halo artifacts at regions of large change in gradients, (e) global contrast stretching, (f) global histogram equalization - both (e) and (f) suffer from saturation artifacts and color blotches.