Toward a Natural Local Color Image Enhancement

Jean-Michel Morel, CMLA ENS Cachan 61 avenue du Président Wilson 94235 Cachan cedex, Paris, France.
Ana Belén Petro, Universitat de les Illes Balears, Crta. de Valldemossa km 7,5 07122 Palma de Mallorca, Spain.
Catalina Sbert, Universitat de les Illes Balears, Crta. de Valldemossa km 7,5 07122 Palma de Mallorca, Spain.

Abstract

There is no unique contrast enhancement method, but the existing methods seem to cover all needs for local or global enhancement. Nevertheless, experimental evidence shows that the two main algorithmic classes (histogram based and high pass filter based) have their own characteristic artifacts. This work intends to show that new and simple alternative algorithms in these classes are possible. The two proposed algorithms introduce bounds on the amount of contrast change, permitting to always deliver natural looking enhanced images. Comparison with six state-of the-art methods illustrate the competitiveness of a fast nonlinear local contrast enhancement method by partial differential equation, and of a very simple histogram modification method.

Introduction

One of the main problems with analogue and digital photography is the discrepancy between the recorded color image and the direct observation scene. Indeed the human visual perception (HVP) made up of eye, retina and several visual cortex areas is a precise and complete mechanism to interpret information from visible light. The HVP has two features which are not owned by cameras, a high local dynamic range and the color constancy. The large tonal difference is not a problem for the human perception, which adjusts locally to brightness changes. The pupil, the iris, and the associated muscles account for the incredible dynamic range of the human eye, which further adapts on a given scene by an active exploration in different gaze directions. A person is therefore able to perceive the vivid color and the details of a scene in both shadows and highlight zones. A fixed camera has no such ability to contrast adjustment and is generally unable to capture all of the tones in the scene. There are not less than four issues explaining why the contrast problem is intrinsically hard: first, an increased time exposure can improve the dynamic range, but increases the risk of motion blur. Second, a short snapshot may avoid blur but decreases the SNR, particularly in dark image regions. Third, the short time exposure could be compensated by increasing the CCD size, but only in detriment of the camera resolution. Last but not least, even if the resulting photograph has a good SNR and no blur, a contrast enhancement remains necessary in the dark regions to mimic the local adaptive human perception.

Contrast enhancement therefore remains one of the most important issues in image processing. Among all image processing techniques, it is the one that has the strongest impact on image quality. Many contrast enhancement techniques have been introduced to improve the contrast of an image but there seems to be no universal method for all applications, simply because the kind of correction depends on the scene.

Our analysis is based on a simple classification of contrast enhancement techniques in histogram methods and frequency domain methods. The histogram methods modify local or global image histograms. Global enhancement methods manipulate the whole image histogram while local enhancement methods use local information, conditioned by fixed or adaptive neighborhood. Their main constraint is that the histogram stretching, local or global, must be controlled to avoid revealing quantization and noise and to avoid squeezing even more some regions.

The most popular global contrast enhancement is the parameterless global histogram equalization (HE) [4]. HE applies a nonlinear transform to the intensity levels so that the cumulative distribution function $F(x)$ becomes closest to linear, i.e. $F(x) \simeq ax$, with a depending on the dimension of the image and the range of values. The local method proposed by Stark in [8] modifies the cumulative distribution function of the histogram to adjust the level of enhancement. Arici et al. [1] propose a global method to modify the histogram by solving an optimization problem. The new histogram is a weighted average of the input histogram and of the desired uniform one. Probably the most sophisticated local histogram adjustment was proposed by Caselles et al. [2], where the contrast change is made conditionally to connected components of level sets. We call this method local contrast enhancement (LHE). Again in this method noise can be exaggeratedly enhanced in some parts.

Other techniques for image enhancement have been inspired by the so called "retinex" theory. One such well acknowledged technique was proposed by Jobson et al. [5], the "multiscale retinex with color restoration" (MSRCR). Jobson et al. compute the logarithm of the responses of several center-surround filters at several scales and combine the results linearly. This method often remarkably enhances image detail, but suffers from several serious drawbacks: by attenuating strongly low pass information, image colors become grayish ; some halo effects near boundaries can also occur. As the GIMP implementation reveals, a good final result can only be obtained by a complex tuning of many parameters for each image, namely the scales and weights of the center-surround filters and several color post-processing parameters. This method based on center-surround contrast can be accounted for as a sophisticated high pass filter. Similarly, the methods enhancing low gradients, enhance local contrast. Such is the Fattal et al. [3] method, providing a contrast enhancement of low dynamic range images.

Although most mentioned methods can actually deliver outstanding results on some images, they all are at risk to excessively alter color or enhance the noise, or depend on a delicate parameter tuning.

In this paper we shall introduce and compare to them two "robust" enhancement methods, one based on a simple global controlled histogram modification, and one of the high pass type manipulating locally the gradient, in the line of the Fattal et al. method, but with contrast control. This method mainly depends on a parameter selecting the dark regions to enhance. While the Fattal et al. method controls the contrast stretch according to gradient properties only, the proposed method increases the local gradients on selected regions and relies on the observation that in most images contrast enhancement is actually only required in dark regions.

Both proposed algorithms tackle most of the issues mentioned above, and appear to always deliver natural-looking images. To that effect, natural bounds are respected to limit the SNR local degradation and color attenuation. The gradient in the image is never allowed to be multiplied or divided by more than a fixed factor for which all experiments show that the ideal value is close to 2.5. A systematic comparison being impossible, due to the lack of benchmark data and to the variety of image contrast issues, all algorithms will be compared on several most difficult cases including images with back light, and the focus will be to demonstrate which kinds of artifacts can be caused (or avoided) by each method.

The plan of this paper is as follows. Section 2 proposes a robust histogram equalization method avoiding artifacts common to this technique. Section 3 is devoted to a local contrast enhancement method which enhances nonlinearly the image gradient in automatically selected regions, and reconstructs the image by integrating the Poisson P.D.E. with a fast Fourier solver. Section 4 compares the two proposed methods with state of the art local and global methods. Section 5 is the conclusion.

Limited slope histogram equalization

Contrast enhancement is always the result of a compromise. Indeed, image colors are constrained to belong to a fixed range. Thus, stretching the contrast in one part may result in squeezing it in another. This also explains why methods cannot be fully automatic. The methods introduced herewith will be no exception, but will be limited to at most two (intuitive) parameters, and actually a main one. The first idea for a global method is to perform a piecewise linear transformation of the intensity levels such that the new cumulative distribution function will be approximately linear, but where the stretching of the range is locally controlled to avoid brutal noise enhancement.

Consider an image of dimension D and with L pixels values in the range [0,255], having a probability density $f(x)$ with cumulative distribution function $F(x)$. We will denote by $F_u(x)$ the linear cumulative distribution function.

We consider a regular partition $0 = y_0 < y_1 < \cdots < y_N = 255$ of the interval $[0,255]$, with $y_k = \frac{255k}{N}$, $k = 0,1,\cdots,N$. For each point y_k of the partition we compute $x_k = F^{-1}(F_u(y_k))$, $k = 1,2,\cdots N-1$ and $x_0 = 0$, $x_N = 255$, we have obtained a new partition of the interval $[0,255]$.

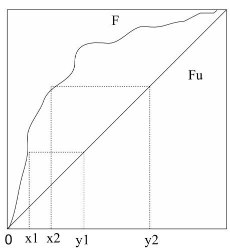

Figure 1. *The distribution of the points x_k and y_k*

For each interval of the new partition $I_k = [x_k, x_{k+1}]$ the algorithm constructs a linear transformation $T_k(x)$ such that T_k transforms the interval $[x_k, x_{k+1}]$ into $[y_k, y_{k+1}]$. Thus,

$$T_k(x) = y_k + m_k(x - x_k) \; k = 0,1,\ldots,N-1, \tag{1}$$

where for each of this linear transformations, the scale factor is $m_k = \frac{y_{k+1}-y_k}{x_{k+1}-x_k}$.

If this scale factor is too small, the corresponding linear transformation can compress the histogram, and consequently could lose too much contrast. On the other hand if the scale factor is too large the transformation stretches the histogram and this can produce a noise amplification, mainly in the dark zones of the image. To avoid this effect the value of the scale factors must be constrained, and we shall define the new scale factor as

$$m_k = \begin{cases} \max(m_k, 0.5) & m_k < 1 \\ \min(m_k, 3) & m_k \geq 1 \end{cases} \tag{2}$$

These new values of m_k modify the partition of the interval [0,255] i.e., the points y_k. Thus the algorithm of the method writes:

- Fix N, yielding the partition points y_k, $k = 0,1,\ldots,N$ of $[0,255]$.
- Compute $x_k = F^{-1}(F_u(y_k))$, $k = 0,1,\ldots,N$.
- For $k = 0,\ldots N-1$ do
  - Compute m_k.
  - If $m_k < 0.5$ then $m_k = 0.5$ and $y_{k+1} = T_k(x_{k+1})$.
  - If $m_k > 3$ then $m_k = 3$ and $y_{k+1} = T_k(x_{k+1})$.
  - Transform by the values $x \in [x_k, x_{k+1}]$ into $y \in [y_k, y_{k+1}]$ by $y = T_k(x)$.
  - $k = k+1$.

This method yielding a piecewise equalization of the image with *limited slope* will be called in the sequel $LSHE$ (limited slope histogram equalization). This simple algorithm, not documented in the literature, turns out to be an excellent competitor for many more sophisticated algorithms.

The number of points N of the partition depends on the image and is the real parameter of the method. When N is large the result is more similar to a histogram equalization algorithm. Figure 2 shows an example of an image which has all of its values concentrated in the dark and middle range. The full histogram equalization (HE) (Figure 2 b)) produces an unnatural effect creating constant zones in the background and in the face of the little rabbit. If we apply $LSHE$ with $N = 3$ (Figure 2 c)) this yields a stretching of the middle range values increasing the contrast while preserving the continuity of the background. Figure 2 d) shows the result with $N = 5$ which is more similar to histogram equalization. The slope parameters instead can be fixed one and for all, since they give sound bounds to contrast lost and SNR degradation.

Local Contrast Enhancement Method

This section introduces a local contrast enhancement based on the enhancement of the gradient of the image in selected regions. The perception of contrast is directly related to the local luminance differences, i.e. the local luminance gradients. Methods which manipulate the gradient need to be integrated into an new image by solving the corresponding Poisson equation. This idea has been used for contrast enhancement and seamless image editing [3], [7]. In Perez et al. [7] the gradient of an image is enhanced in some region which is selected manually. The corresponding Poisson equations is solved by a multigrid method.

Here we propose a far more automatic and complete method where the selection of the regions to be enhanced is specified by a single parameter, and the regions are allowed a complex topology without additional computational cost.

The idea is to define a new gradient field such that in the low contrast regions a enhancement function on the gradient is applied. Given an image f with domain R, we denote by $\Omega \subset R$ the selected regions on R, and define a vector field $\mathbf{V}$ as

$$\mathbf{V} = \begin{cases} \varphi(\nabla f) & \text{over } \Omega \\ \nabla f & \text{otherwise.} \end{cases} \tag{3}$$

©2012 Society for Imaging Science and Technology

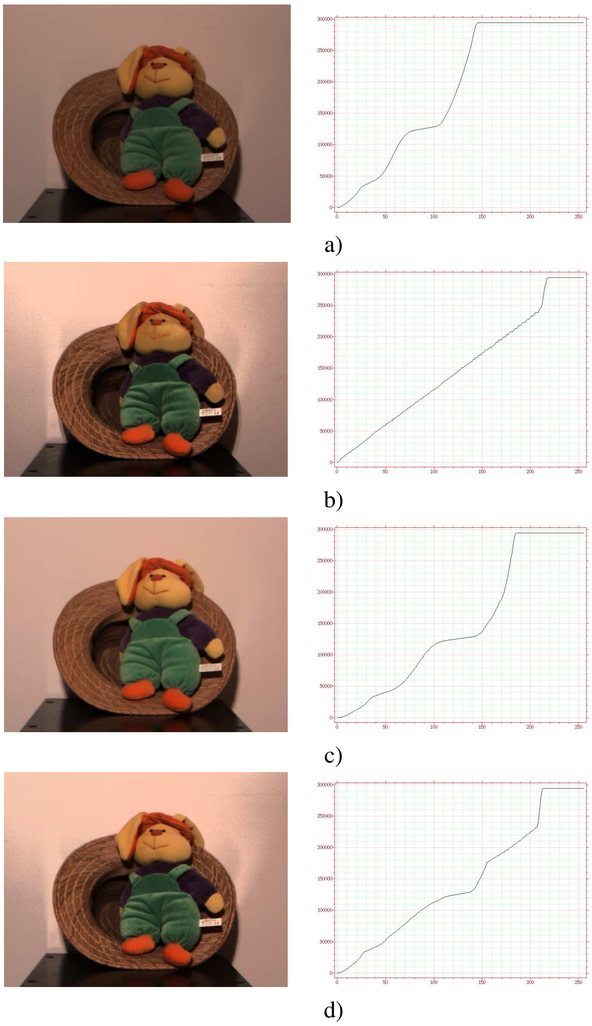

Figure 2. a) Original image and the cumulative distribution function. b) The result of HE and its CDF. c) The result of $LSHE$ with $N=3$. d) The result of $LSHE$ with $N=5$.

Given this vector field $\mathbf{V}$, the problem is to find the image u whose gradient field is the closest, in L_2-norm, to the prescribed vector field $\mathbf{V}$ i.e. the problem is to solve the following optimization problem,

$$\min_u \int_R |\nabla u - \mathbf{V}|^2. \tag{4}$$

The minimizer is determined by the Euler-Lagrange equation with homogeneous Neumann boundary condition

$$\Delta u = \operatorname{div} \mathbf{V}, \quad \text{over } R, \quad \frac{\partial u}{\partial \mathbf{n}} = 0 \quad \text{over } \partial R, \tag{5}$$

where $\mathbf{n}$ is the direction orthogonal to the boundary. This mathematical problem has a unique solution up to an additive constant. The final solution is obtained by fitting the min and max of the solution to the maximal possible interval (usually $[0,255]$). An analysis of a wide range of ill-contrasted images shows that an overwhelming part of the contrast problems is located in the darker image regions. Thus, we found that the best choice for Ω is to take an image lower level set,

$$\Omega = \{x \in R : \ f(x) < T\}.$$

Figure 3. Top: Original image. Middle: Result of local gradient enhancement in the dark region with $\varphi(x) = 2.5x$, in which case the algorithm is equivalent to LHE. Bottom: Result with $\varphi(x) = sign(x)|x|^{0.8}$, with sharper detail.

Surprisingly, and although this parameter should be left to the user, $T = 50$ gave good results in all experiments, probably because most computer screens lack contrast below this value.

For the enhancement function φ, the simplest possibility is a linear function, which amounts to simply scaling the gradient of the dark regions. In that case it can be proved that the algorithm proposed below is a particular instance of LHE [2]. Nevertheless, concave power functions give a still better result because they do not unnecessarily enhance high gradients in dark regions. Figure 3 shows an example of this method where we have used $\varphi(x) = 2.5x$ (Figure 3 middle) and $\varphi(x) = sign(x)|x|^{0.8}$ (Figure 3 bottom).

In conformity with Fourier's original method, the Fourier transform can be used to solve the Poisson equation (5), which is faster than the multigrid method proposed in [7] and [3] and gives an exact solution. The Neumann boundary condition is implicitly imposed by extending the original image symmetrically across its sides, so that the extended image, which is four times bigger, becomes symmetric and periodic.

In short, the strategy for solving (5) by Fourier technique is

- Quadruplicate by symmetry the discrete domain and $\mathbf{V}$;
- Compute the discrete Fourier transforms of V_1 and V_2;
- Compute the discrete Fourier transform of the solution $\widehat{u}_{mn}$ as $\widehat{u}_{mn} = \frac{\frac{2\pi i m}{J}\widehat{V}_{1mn} + \frac{2\pi i n}{L}\widehat{V}_{2mn}}{\left(\frac{2\pi m}{J}\right)^2 + \left(\frac{2\pi n}{L}\right)^2}$;
- Obtain the samples u_{jl} of the solution by the inverse discrete Fourier transform;

- Restrict them to the initial domain.

Remark: The Fourier transform comes with some overhead such as padding to the nearest power of 2, but with smart Fourier libraries, like fftw, having integer values that are not powers of 2 is no more a complexity issue. Products of small factors are most efficient, but and $O(n \log n)$ algorithm is used even for prime sizes. Likewise, we mentioned for pedagogic reasons the quadruplication. In fact, again with smart library like fftw the quadruplication is implicit and performed directly as a cosine transform.

Results and Comparison

Unfortunately there is no objective criterion to compare contrast enhancement techniques. For most metrics evaluating the contrast in an image, HE achieves usually the best results because it maximizes the image entropy. Nevertheless, the contrast enhancement methods in the literature suffer from some artifacts that make the image unnatural. In this paper we will define some of these artifacts:

Quantization noise: reveals level lines as shock lines, creating unnatural edges in the images, typical of histogram methods;
Halo: luminance oscillation near the strong edges, typical of the high pass filters;
False colors: caused by excessive changes in color dynamics;
Loss of hue: colors become grayish, typical of high pass filters;
Saturation: loss of contrast in bright regions, typical of histogram methods;
Excessive texture enhancement: typical of high pass filter. usually combined with a loss of color and global contrast this results in unnatural images with fantastic detail.

This section presents a visual comparison of the proposed two methods with six state of the art methods illustrating how the natural bounds proposed in these methods control the artifacts caused by others. Table 1 classifies by artifacts the various methods. Of course a method cannot be adopted just because it does not create artifacts, the identity being the best method for this criterion alone. Thus a visual comparison is necessary to check that the enhancement goal has been attained.

The eight compared methods are the limited slope histogram equalization LSHE, the local contrast adjustment LCA, both presented in the previous sections, the classic contrast enhancement histogram equalization HE [4], the simplest color balance algorithm SCB [6], the adaptive contrast enhancement ACE [8], the local histogram equalization LHE [2], the multiscale “retinex” with color restoration MSRCR [5] and the Fattal et al. gradient-domain high dynamic range compression HDRC [3].

The simplest way to stretch or to contract the histogram of an image is defined by

$$T(l) = \frac{s_{max} - s_{min}}{l_{max} - l_{min}}(l - l_{min}) + s_{min} \tag{6}$$

where l_{max} and l_{min} are the maximum and minimum values of the intensity levels in the image, s_{max} and s_{min} are the desired maximum and minimum values in the histogram. If the factor $\frac{s_{max}-s_{min}}{l_{max}-l_{min}}$ is smaller than 1 this transformation contracts the histogram. If it is larger than 1 the transformation stretches the histogram, and if $s_{max} = 255$ and $s_{min} = 0$ the transformation stretches, as much as it can, the histogram. However, many images contain a few aberrant pixels that already occupy the 0 and 255 values. Thus, an often spectacular image color improvement is obtained by “clipping” a small percentage of the pixels with the highest values to 255 and a small percentage of the pixels with the lowest values to 0, before applying the affine transform (6), the percentage of saturated pixels must be as small as possible. This algorithm called “simplest color balance”, SCB, is proposed in [6], and can be applied to the intensity histogram or to each R, G, B channel histogram, depending on the application.

In the comparison below, results of SCB come from the on line article in *http://www.ipol.im*. The results of HE, ACE and MSRCR were obtained with corresponding GIMP plug-ins. The parameters of ACE and MSRCR we have used are the plug-in default parameters. The results of HDRC were obtained from the web page *http://www.cs.huji.ac.il/ ∼ danix/hdr/enhancement.html*. The authors of LHE [2] provided us the code for the experiments. The proposed methods and LHE method were applied to the luminance component. The output color image was obtained by multiplying each color channel of the input image with the ratio of their output and input luminance values to preserve the hue.

The compared methods in divided in two coherent groups: a) histogram modification based methods (HE, SCB, ACE, LHE and LSHE); b) high pass filter methods (MSRCR, HDRC and LCA).

Figure 4 shows the original test images and their corresponding contrast enhanced with the algorithms based on the modification of the histogram: SCB, HE, ACE , LHE and LSHE. The simplest color balance (Figure 4 top middle) does not present any important changes with respect to the original, since the range of values of the original image is $[0,255]$. HE (Figure 4 top right) is probably the most contrasted result but it does not mean that the resulting image is better in terms of visual quality. The image has an unnatural look with excessive texture enhancement and false colors. ACE (Figure 4 bottom left) produces a better contrasted image but there are some halos on the sky and on the lake and presents some quantization noise. LHE (Figure 4 bottom middle) and the proposed method LSHE (Figure 4 bottom right) produce similar results. Nevertheless a serious difference is observed in the lake and in the sky, where LHE produces an unnatural sharp transition of color in the lake revealing quantization, while LSHE ($N = 4$) produces a smoother transition. Table 1 summarizes the artifacts creates by each method.

Figure 5 shows the results of high pass filter methods, HDRC, MSRCR and the proposed method LCA. The HDRC result, Figure 5 top right, is probably the best result, but Fattal et al. work in a multiscale framework, while the proposes method works in a single scale only, therefore it is much simpler and the results are comparable. The swan in the MSRCR, Figure 5 bottom left, loses the texture and shadows of feathers, the swan is saturated to white, and the blue flowers have lost the color.

Figure 6 shows a comparison of all the methods, except HDRC. The SCB result is exactly the same as the original image, because the range covers the interval $[0,255]$. HE produces an excessive texture enhancement on the rock and false colors. ACE produces quantization noise and excessive texture enhancement. LHE produces quantization noise mainly on the sky. MSRCR creates saturation in the sky and the image becomes grayish. Like HDRC but simpler, LSHE and LCA produce more natural images and do not create artifacts. See Table 1.

Conclusion

This paper has specified two methods for contrast enhancement, one global with low computational cost and another far

Account of the artifacts caused by the eight compared methods

	HE	ACE	LHE	MSRCR	HDRC	LSHE	LCA	SCB
Quantization noise	X	X	X					
Halo		X		X				
False colors	X							X
Loss of hue		X	X					
Saturation				X				X
Excessive texture	X	X		X				

Figure 4. *Comparison of the histogram based methods. Top left: Original image. Top middle: SCB result. Top right: HE result. Bottom left: ACE result. Bottom middle: LHE result. Bottom right: 4 points $LSHE$ result.*

more sophisticated based on Poisson editing. In both methods the principal novelty is a careful control of noise enhancement and the preservation of hue. We compared both methods to six other state of the art methods. The outcome of such a comparison being necessarily subjective, we picked characteristic difficult examples. For the histogram methods, LSHE is extremely simple and guarantees good results in most cases, while not being able to tackle images with dark regions, where a local enhancement is necessary. In all cases, the local contrast enhancement (LCA) seems to give the best attainable result, while creating no artifact, like HDRC. The obvious drawback of this last algorithm is the dependence on two necessary parameters, one selecting the extent of dark regions to enhance, and one selecting the parameter to amplify the gradient. But even these parameters are intuitive and were fixed to the same values for all images. Both methods actually share a manipulation of image gradient, which seem to avoid the artifacts of more global high-pass filter methods like retinex (MSRC).

References

[1] T. Arici, S. Dikbas, and Y. Altunbasak. A histogram modification framework and its application for image contrast enhancement. *IEEE Trans. on Image Processing*, 18(9):1921–1935, 2009.

[2] V. Caselles, J.L. Lisani, J.M. Morel, and G. Sapiro. Shape preserving local histogram modification. *IEEE Transactions on Image Processing*, 8(2):220–230, 1999.

[3] R. Fattal, D. Lischinski, and M. Werman. Gradient domain high dynamic range compression. *ACM Transactions on graphics*, 21(3):242–256, 2002.

[4] R. C. Gonzalez and R. E. Woods. *Digital Image Processing. Second edition*. Prentice Hall, 2002.

[5] D. J. Jobson, Z. Rahman, and G. A. Woodell. A multiscale retinex for bridging the gap between color images and the human observations of scenes. *IEEE Transactions on Image Processing*, 6(7):965–976, 1997.

[6] N. Limare, J. L. Lisani, J.M. Morel, A. B. Petro, and C. Sbert. Simplest color balance. *IPOL, Image Processing on Line, http://www.ipol.im*, 2011.

[7] P. Perez, M. Gangnet, and A. Blake. Poisson image editing. *ACM Transactions on Graphics*, 22(3):313 –318, 2003.

[8] L. Alex Stark. Adaptive image contrast enhancement using generalizations of histogram equalization. *IEEE Trans. on Image Processing*, 9(5):889–896, 2000.

Author Biography

Jean-Michel Morel (PhD 1980) is Professor of Applied Mathematics at the Ecole Normale Supérieure de Cachan since 1997. His research is focused since 1990 on the mathematical analysis of image analysis and processing. He has coauthored with S. Solimini a book on "Variational Methods in Image Segmentation" (Birkhäuser 1994). He has also co-authored with Agnès Desolneux and Lionel Moisan From Gestalt theory to image analysis: a probabilistic approach (Springer, 2008), and is also coauthor of A Theory of shape identification (Springer LNM, 2008).

Ana Belén Petro (PhD 2006) is a Full-time Lecturer at the Universitat de les Illes Balears. Her research is focused on the color image analysis and processing.

Catalina Sbert (PhD 1995) is Associate Professor at the Universitat de les Illes Balears since 1999. Her research is focused on the mathematical image analysis, particularly she is interested on the partial differential equations applied to the image processing.

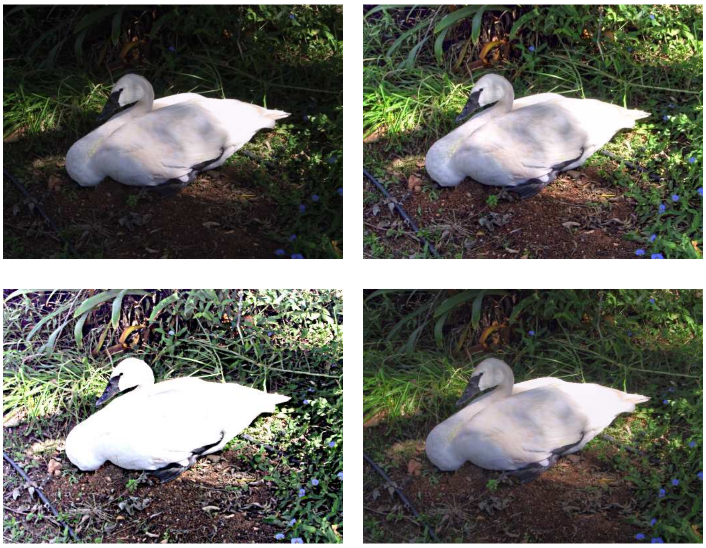

Figure 5. *Comparison of high pass filters. Top left: original image. Top right: HDRC result. Bottom left: MSRCR result. Bottom right: LCA result.*

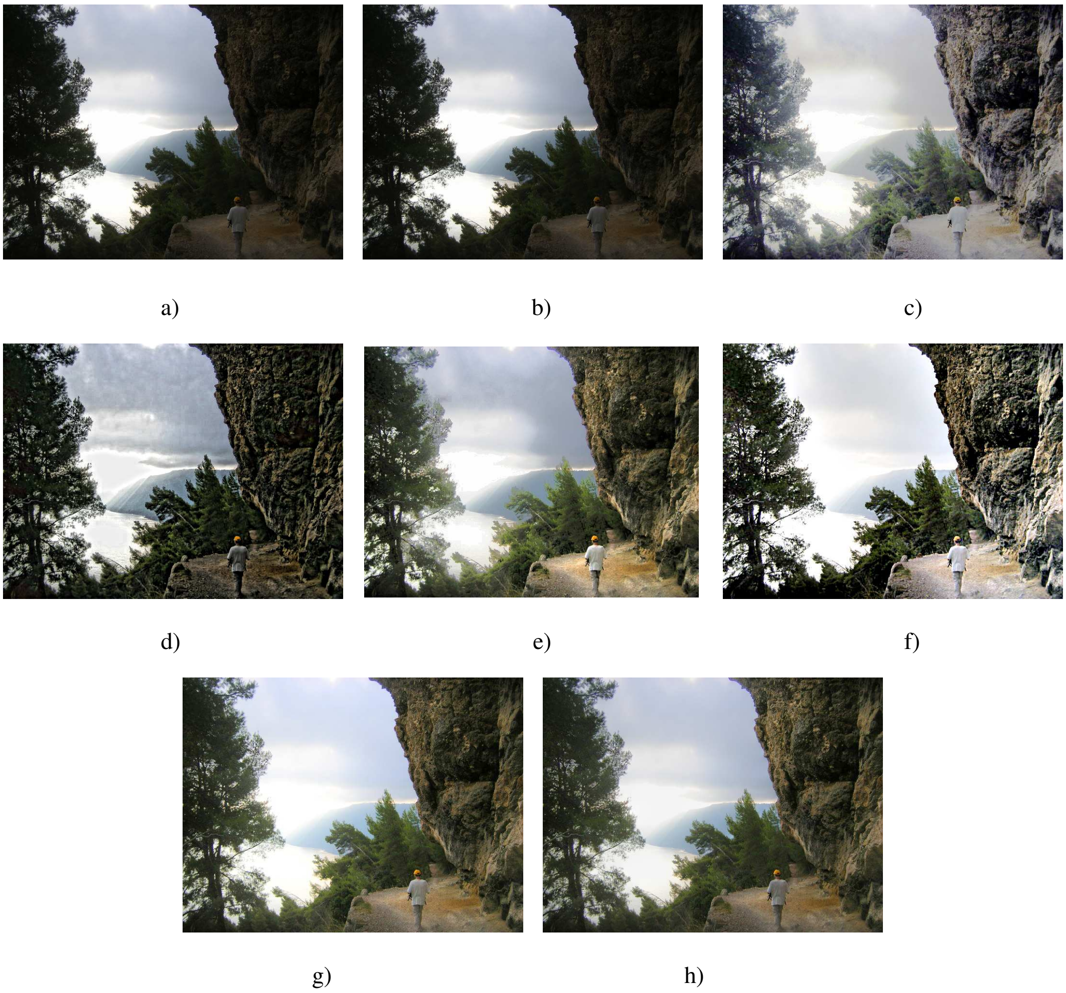

Figure 6. *a) Original image. b) SCB result. c) HE result. d) ACE result. e) LHE result. f) MSRCR result. g) 2 points $LSHE$ result. h) LCA result.*

©2012 Society for Imaging Science and Technology