

IMPROVED ILLUMINATION CORRECTION THAT PRESERVES MEDIUM-SIZED OBJECTS

Anders Hast¹ and Andrea Marchetti²

¹*Department of Information Technology, Uppsala University, Uppsala, Sweden*

²*Consiglio Nazionale delle Ricerche, Institute of Informatics and Telematics, Pisa, Italy*

Abstract. Illumination correction is a method used for removing the influence of light coming from the environment and of other distorting factors in the image capturing process. An algorithm based on the luminance mapping is proposed that can be used to remove low frequency variations in the intensity, and to increase the contrast in low contrast areas when necessary. Moreover, the algorithm can be employed to preserve the intensity of medium-sized objects with different intensity or colour than their surroundings, which otherwise would tend to be washed out. Furthermore, examples are given showing how the method can be used for both greyscale images and colour photos.

Key words: illumination correction, luminance mapping, image stitching, image mosaicing, vignetting.

1. Introduction

Varying illumination in captured images is a common problem in many different areas, such as Photography, Microscopy and Aerial imaging. Hence different names are used for both the problem and the remedies across different disciplines. Vignetting (Yu, 2004 [19]; Zheng et al, 2006 [24]) occurs due to different mechanisms (Goldman and Chen, 2005 [20]) that cause brightness falloff away from the image centre, and is prevalent in photography. The mechanisms are primarily: *mechanical-, optical- and natural vignetting* (Sidney 2002 [12]). The first one is due to the incoming light being blocked by external objects, such as filters or lens hoods. Optical vignetting, on the other hand, is due to the dimensions of the lens, and can be remedied by closing down the aperture. The full aperture will yield gradual darkening towards the corners, and thus a brighter spot in the middle. Natural vignetting, or natural illumination falloff, is proportional to the fourth power of the incoming light angle cosine. The remedy is to use a gradual grey filter or post-processing, like the method proposed herein, which we will refer to as *Illumination Correction*.

Non-uniform illumination, or even dirty lenses and dust (Young, 2000 [9]), cause problems in Microscopy (Leong et al, 2003 [14]), and the remedying procedure is often called illumination or shading correction. In Magnetic Resonance Imaging (MRI), varying shade is known as RF-inhomogeneity or bias (Agus et al, 2007 [25]; Ardizzone et al, 2007 [26]). In Face Recognition, varying illumination is a challenging problem (Guo

et al, 2005 [23]; Zhu et al, 2003 [17]) and the remedying method is, besides illumination correction, also called illumination normalization. The illumination itself is often referred to as background light. In Mammography, contrast-limited adaptive histogram equalization (CLAHE) has been used (Puff et al, 1994 [5]). In aerial and satellite photos, physical lighting also affects relief presentation (Rocchini et al, 2005 [22]; Tzelepis et al, 2003 [16]). However, it is also a problem for image stitching and mosaicing of panoramas, where the focus often is put on making the transition from one image to the other as smooth as possible (Burt and Adelson, 1983 [2]; Levin et al, 2004 [18]; Pérez et al 2003 [15]).

In this paper we will focus on a method that was previously developed for *retrospective illumination correction of greyscale historical aerial photos* (Hast and Marchetti, 2011 [30]). The method can be used both for minimizing the change in intensity over the image and for increasing the contrast in low contrast areas. This is important, as the contrast is generally lower in high and low intensity areas. We also present an important enhancement of the method, aimed at preserving the original intensity of medium-sized objects having intensity different than that of their surroundings. Here by *medium-sized objects* we mean objects that are much smaller than uneven illumination varying across the image, but still much larger than small details in the image. Preservation of their intensity would otherwise be impossible, as the goal of any illumination correction method is to minimize the changes in intensity across the image, which will affect such objects. Furthermore, we show how the new method can also be used for colour images, and that it will help preserve colours better.

2. Previous Research

There are many methods for illumination correction, each of them tailored to the specific field where it is used, as mentioned in the introduction. There are many different ways to handle the problem of varying illumination. One of the first approaches published was to use a homomorphic filter (Oppenheim et al, 1968 [1]), which takes advantage of the fact that the intensity in the image can be divided into two components: the illumination of an object and the amount of light reflected by that object (Pajarez et al, 2005 [21]). A homomorphic filter uses Fourier transform to remove from the image illumination, which is characterized by low spatial variations, and to keep reflectance, which is characterized by high frequency changes. Hence multiplication needs to be converted to addition, which can be done by taking the logarithm of the functions, as Fourier transform only can be used when the noise is an additive term. The main disadvantage of this approach is that it requires the use of image size padding when the image size is not a power of 2, which can lead to distortions in the boundary regions (ERDAS, 2010 [29]).

Another way to handle the problem (Zhu et al, 2003 [17]) is to model the lighting change as a local affine transformation (AT) of the pixel value, where the low spatial

variations caused can be removed using a multi-resolution low pass filter that estimates the background lighting (Zhu et al, 2003 [17]). A similar approach can be adopted using a Gaussian filter (Leong et al, 2004 [14]). Others handle the background light using a reference picture acquired in a controlled environment with known lighting conditions and a white object (Yu, 2004 [19]). This is, of course, not always possible, especially in retrospect. Yucong et al (Yucung et al 2005 [23]) use histogram equalization combined with ICR (Illumination Compensation based on Multiple Regression Model) (Ko et al, 2002 [11]), which aims at finding the plane that best fits the intensity distribution in the image using a multiple regression model. Then this plane is used to remove illumination of face images.

Local Range Modification (LRM) (Fahnestock and Schowengerdt, 1983 [3]) finds interpolated minimum and maximum pixel values in a neighbourhood (contextual region), and stretches them to the desired range. Similarly, the contrast-limited adaptive histogram equalization method (CLAHE) (Pizer et al, 1987 [4]) stretches the histogram in its contextual region. This improves the image in such way that the contrast is enhanced. In microscopy (Yuong et al, 2000 [9]) make use of both homomorphic filters and morphological filtering, which yields an estimate of the background illumination employed to correct images. Others have developed methods for finding the *vignetting function* from a single image (Zheng et al, 2006 [24]) or from multiple images, with or without the response curve (Goldman and Chen, 2005 [20]). Wavelet-based approaches for image and contrast enhancement are also popular (Reeves and Jernigan, 1997 [8]), and can be used to enhance contrast in digital mammography (Laine et al, 1995 [7]). Yet others make use of a combination of methods, such as wavelets and the homomorphic filter (Ashiba et al, 2008 [27]; Yoon and Ro, 2002 [13]).

3. The Proposed Method

A part of the proposed method was recently presented as a poster (Hast and Marchetti, 2011 [30]). Here we present the novel idea of preserving medium-sized areas with different intensity, and show how the new method can be used for colour images.

3.1. The Main Idea

The main idea of the previous paper was to use the luminance mapping (Hertzmann et al, 2001 [10]), which is usually employed in colour correction (Vrhel and Trussel, 1994 [6]). Here a linear map is applied that matches the means and deviations of luminance distributions, i.e. the intensity. If $A(x, y)$ is the luminance of a pixel in image A , then it is remapped using the distribution of image B :

$$A(x, y) = \frac{\sigma_B}{\sigma_A} (A(x, y) - \mu_A) + \mu_b . \quad (1)$$




Fig. 1. The image on the left contains an obvious shading artefact, visible as intensity falloff away from the centre. The diagram on the right shows a cut section taken in the middle depicted in blue, and the local mean – in red.

where μ_A and μ_B are the mean luminances, and σ_A and σ_B are the deviations of the luminances, both taken with respect to the luminance distributions in A and B , respectively. The result of such remapping, if applied to all three colour channels, is that image A will have colours looking similar to those in image B .

Let us take as an example the image of the checkerboard on the left of Fig. 1, which has an obvious intensity variation. Looking at the horizontal cut section taken from the centre and depicted on the right, we can see the intensity variation over the cut. The *local mean* is the red curve that runs along the centre of the curve. If the *local mean* is translated to the *global mean*, i.e. transforming the curve to a line, then a major portion of the varying light is removed.

A *low-pass filter* is used to remove an estimate of the illumination that contains the lowest frequencies in the image, and thus an approximation of the *local mean* μ_A . This is done by using a Gaussian filter (3) modified in such a way that it also compensates for the borders. The trick of the trade is to divide the result by a likewise smoothed image that before the smoothing contained the intensity 1.0 everywhere. The smoothed image will hence contain 1.0 everywhere except at the borders, where the values will be lower – and this will compensate for the corresponding effect in the input image. Nevertheless, the mask must be quite large, up to the input image size.

Since the illumination is computed using very large Gaussian masks, this approach becomes impractical even with modern computers for very large photos. In the case of historical aerial photos, we use images of sizes up to 5000×6000 pixels, and the mask can be just as large. Nonetheless, it is still possible to compute the illumination using a downscaled version \mathbf{h} of the captured image \mathbf{H} , which is therefore blurred much faster,

and then is resized to its original size. Since the result will be a heavily blurred image containing smooth and slow variations, an interpolation of such image will affect the result very little. The local mean is hence computed as a Gaussian filtered image using the convolution:

$$\mu_a(x, y) = G * \mathbf{h}(x, y) \quad (2)$$

where

$$G(x, y) = \frac{1}{2\pi\sigma^2} e^{-(x^2+y^2)/2\sigma^2} . \quad (3)$$

The width of the distribution used to delimit the low frequencies is determined by σ . The larger the value of σ , the greater the smoothing effect. The resulting blurred image will be very close to what the full scale Gaussian filter would produce, and will not be detrimental to the final quality, as long as the downscaling is reasonable. However, it should be mentioned that in Figs. 1 and 2 the scale factor of 1 was used, i.e. there was no downscaling at all.

In the next step, the local deviation σ_A is computed as the Gaussian blurred absolute value of the difference between the downscaled image \mathbf{h} and μ_a :

$$\sigma_a(x, y) = G * |\mathbf{h}(x, y) - \mu_a(x, y)| . \quad (4)$$

The local deviation σ_A is computed by upscaling σ_a bilinearly to the full size. The illumination-corrected image \tilde{H} is computed as:

$$\tilde{H}(x, y) = \frac{\sigma}{\sigma_A(\mathbf{x}, \mathbf{y})} (H(x, y) - \mu_A(x, y)) + \mu . \quad (5)$$

Here σ is computed as the mean of σ_a , i.e. the mean variance in the whole downscaled image, and μ is the mean of μ_a , i.e. the overall mean intensity of that image. This equation is used in three steps with the same size of the Gaussian kernel for computing Eqs. (4) and (5). The result is shown in Fig. 2. The proposed contrast enhancement using the luminance mapping yields a straightened curve, where the amplitude is almost the same along the whole curve and the shading artefact is removed from the image.

The proposed method yields satisfactory results for images containing textures with small variations, even if they are quite drastic, as the checkerboard in Figs. 1 and 2. However, if the image contains medium-sized objects with different intensity than their surroundings as in Fig. 3 (left), then the result obtained using the method as explained so far will be quite disappointing, as shown to the right. The black and white squares (medium-sized objects) tend to get the same grey intensity as their surroundings, and give the impression of being washed out.

Fig. 4 shows on the left a cut section of the same image as that on the left in Fig. 3. A Gaussian filter is used to compute the mean, depicted in red, using different σ 's in the top and bottom rows. The corresponding results of the compensation are shown




Fig. 2. The image from Fig. 1 and the cut section after applying the proposed algorithm. The low frequency has been removed, and the contrast – enhanced.




Fig. 3. The image on the left contains, except for vignetting, also one black and one white object. As seen to the right, these will be washed out, since with the original version of the proposed algorithm they tend to get the same intensity as the rest of the image.

on the right. Obviously, the result shown in the top row has lower variation along the whole curve. However, in the lower row the black object is preserved much better, even if the curve is not so straight elsewhere as in the top row, i.e. some illumination remains. Hence a solution is needed that gives us the best of both cases.

The new idea is to incorporate a *band pass filter* in the original approach, which temporarily diminishes the objects having different intensity by simply removing them from the original image. Hence, what is done is actually band rejection, following which a low pass filter is used to remove the varying intensity in the image. Finally, the objects are reinserted by simply adding the band pass-filtered image. This process is also repeated using different mask sizes for the low pass filter, like in the original approach. The band pass filter will therefore be chosen depending on the size of the objects that are to be preserved.

Fig. 5 shows how the image in Fig. 3 (left) is band pass-filtered (top left), and the original image is suppressed (top right) by the band pass image. The lower left picture shows the result after one iteration, where the low pass-filtered image is enhanced by the band pass image. Clearly, it has lower intensity variation, and the black and white squares are better preserved. Nevertheless, as there is some small variation across the




Fig. 4. On the left there is a cut section of the image that cuts through the black square. The average is depicted in red, using Gaussian filters with different values of σ . The corrected curves are shown on the right. The top curve has a lower variation over its whole length. The bottom curve preserves the black square, but is more bent elsewhere compared to the top one.

band pass-filtered image, the whole process needs to be repeated, just like in the original method. The result after three iterations is shown on the upper right. The varying illumination has been removed, and the objects are preserved.

In our tests we used a band pass filter known as a difference of a Gaussians filter (DOG) (Russ 2011 [31]). It is constructed as the difference of two Gaussian low pass filters with different values of σ . Hence it yields an image that contains medium-size objects, but also removes a major part of low variations in the image, as well as high frequency details. Together with increasing the contrast in low-contrast regions, the algorithm will also decrease the contrast in the areas where it is already high. Hence it is often necessary to avoid contrast enhancement in those areas. This can be easily done by not allowing σ_A to exceed some predefined percentage of σ_B . In this manner, we can adjust the image to obtain the desired result.

3.2. Repeating and Halving

It was found earlier that repeating the algorithm a couple of times using different σ 's gave a more pleasing result, and this is still true for our new approach. At first, such a method was applied only because it helped preserve medium-sized objects of different intensities, but it turned out to have a rather small impact on such objects. What is more important is that it helps remove low frequency variations in a better way, especially




Fig. 5. The result of the band pass filter is shown on the upper left (scaled to the $[0..1]$ range), and the result of its adding to the input image – on the upper right. The lower left picture shows the result of one step of the previously explained illumination correction, involving the removal of the band pass-filtered image. Clearly, the fields are not suppressed and the shading is corrected; however, more cycles are needed to give the desired result, which is shown on the lower right.

when increasing the contrast, than using a mask of the same size one or several times, as shown in Fig. 6. There is nothing that prevents us from repeating the method fewer or more times than three. However, empirical tests have established that three times are often good enough, and that more iterations do not improve the result much. In any case, this approach needs to be used when preserving medium-size objects, as the band pass-filtered image contains some small variations. Repeating the process gradually eliminates these variations. It has been noted that halving of the mask gives both less distortion and fewer problems at the image edges. Moreover, it has been found that doubling of the mask size from size a to b gives an almost identical result as halving the size from b to a . In the method presented here we have chosen to halve the mask size in each step.

3.3. Scaling

The original approach proposed that both the local mean and the deviation be obtained by a Gaussian filter, where both can be computed from downscaled versions of the original image, and then upsampled by bilinear (or bicubic) interpolation. This will decrease computation time substantially, as Gaussian convolution is computationally expensive; however, it will not affect the result for moderate downscales.




Fig. 6. A close-up of the curves obtained using different techniques for removing the low frequencies and increasing the contrast. The blue curve is the proposed method, which includes repeating the process three times while halving the mask size. The green curve is the result of using the same mask size three times. It is generally more bent at the tops and bottoms, where it should be straight. The red curve is the result of one iteration only, and it is straighter at the tops and bottoms, but decays too rapidly at the image edges.




Fig. 7. A comparison of a cut section in the blurred image, using a scaled image (200×200 , green curve) for Gaussian convolution and the full image (2048×204 , blue curve). Usually, the difference is negligible for reasonable scales. Left: full graph; right: close-up of its central part.

As shown in Fig. 7, the difference between blurring a smaller image using Gaussian convolution and then interpolating it to full scale compared to blurring the full scale image directly is negligible. The reason is that the blurred image becomes smooth, and interpolation of a smooth curve still yields a smooth curve as long as the scale is not too large. Nonetheless, Gaussian convolution is time-consuming and, for instance, the use of the full size image took ten times more time to compute in our implementation of the algorithm than scaling each side by four (13.0 s compared to 1.3 s for a 800×800




Fig. 8. A cut section of the image that intersects the black square. The red curve is the original image with a clearly visible illumination artefact. The blue is the illumination-corrected image, where the curve is much more straight.

image). Hence the proposed method of working with downscaled images is an important improvement when it comes to computational cost, while still maintaining the quality.

The result of using the proposed algorithm for the artificial image in Fig. 3 (left) is shown in Fig. 8. The red curve is a cut plane of the image in Fig. 3 (left), and the blue curve is a cut plane of the resulting image (Fig. 5, bottom right). The blue curve is straighter, which means that illumination has been removed, while the object is still preserved.

4. Results

Let us examine some results and possibilities of using the proposed method for illumination correction. Fig. 9 shows a historical photo of the river Arno and the west part of Pisa taken during WW-II.

In the top row there are some cropped original photos with different degrees of uneven illumination. The photos in the middle row have been processed with the original method. Both the large, bright field in the leftmost image and the dark field at the top of the rightmost image, as well as the river, have become greyish and are washed out. One can also note a ripple effect on a thin line outside the borders of the field and the river, which is a bit darker outside the field and a bit lighter outside the river. In the bottom row, the new approach is used. As a result, both the field and the river now have more accurate intensity, and the ripple effect is no longer visible.

4.1. Colour Images

Figure 10 shows some results of applying the proposed method to colour images. The leftmost image in the top row shows a wall, and the proposed method was applied to the V component of the image transformed to the HSV space (Sonka, 2008 [28]). This allowed for working solely on the image intensity, regardless of the colour. The said




Fig. 9. ©MiBAC-ICCD, Aerofototeca Nazionale, fondo RAF. A series of historical photos of Arno close to Pisa taken during WW-II in the top row, showing different degrees of uneven illumination. In the middle row, the original method is used to correct the lighting. Both the large fields and the river are washed out, and their borders also show a ripple effect. In the bottom row, the new approach is used. Both the fields and the river have more accurate intensities, and the ripple effect is no longer visible.




Fig. 10. ©Anders Hast. Top row, from left to right: Wall in Pisa, its uneven illumination, wall corrected using the original method (HSV, V only), band pass image, corrected using the new method (HSV, V only). Middle row, from left to right: Cake, corrected using the new method (HSV, V only), corrected using the new method (RGB). Bottom row, from left to right: wall with small intensity variation, corrected with the original method (RGB), with the new method (RGB), and again with the new method (HSV, V only).

image contains very little variation in illumination, as seen in the next image to the right. The subsequent image to the right shows the result of using the original method, and we can see that the whiter parts of the stone texture have become grey and dull. With the new method, these whiter parts become separated (the second image to the right). They are used for compensation according to the proposed method, and the result is the rightmost image. The middle row shows a cake in a rather large variation of illumination. Since the upper part is close to black, the conversion between RGB and HSV yields an undetermined colour, and the upper corners in the image are not represented correctly. The image to the right is processed on all three colour channels (RGB) directly. Now the upper corners of the image are represented correctly, and the result is better.

In the bottom row there is a wall with low intensity variations (left). When working on RGB with the old method, we obtain incorrect colours, as can be seen in the next



Fig. 11. ©Anders Hast. Left: Sunset on the beach. Right: image corrected using the new method (RGB). The beach is now more visible, while the sky has become darker.

image to the right. The dark iron bar, as well as the white part of a brick, both marked by blue circles, tend to become reddish. In the first image to the right the new method has been used for RGB, while in the second image to the right the new method has been employed to process the V part of HSV. The results for the latter two are very similar, and the colours in the mentioned critical parts are obviously better represented. The choice of the method is determined by the amount of colour information present in the image. If there are parts which are close to either black or white depending on uneven illumination, then working on all three colour channels (RGB) is necessary; otherwise the alternative yielding the best results can choose.

The new method can also be used to attenuate parts of an image that are very dark, and where the details are barely visible. Hence the proposed method need not be used just for normalizing the illumination in the whole picture, and can be employed for processing images to obtain other effects. Figure 11 shows (on the left) a photo of a sunset, where the chairs on the beach are barely visible. To the right there is the result of processing each colour channel separately. Now the chairs are clearly visible, and the method allows for enhancing the contrast so that the image seems to have been taken in much stronger light than there actually was. The drawback is that the sky becomes darker, as the algorithm tries to equalize the differences in illumination. This can easily be handled by interpolating between the images to obtain a desirable result.

5. Discussion

When using the proposed method in practice, the user must provide some parameters. Scaling was briefly discussed in 2.3, and can make the whole process a lot faster. Nonetheless, the scale factor cannot be too large. For example, the scale factor of 10 is almost 60 times faster (2.4 s in matlab) in our implementation than the scale factor

of 1 (151 s in matlab) for the rightmost photo in Fig. 9. However, the mean difference between the images is less than 0.08%, and is not noticeable in their visual inspection. Nevertheless, we propose that the scale parameter be set as low as possible while still giving acceptable total computational time.

The number of iterations was discussed in Section 2.2, and the value of 3 was used for the images we have tested, even if the difference sometimes was very small when using just one or two iterations. Nonetheless, the tests indicate that four or more iterations do not give any visible improvement, and therefore we have used the default value of 3 in our tests.

Other parameters are the σ 's for the low and band pass filters discussed in Section 2.1. The first must be set depending on how much of the image is occupied by uneven light. In our experience, this is quite similar for most of the images in a set of aerial images taken during one flight, so a default parameter can be chosen for the first image, and then reused for the others. The second σ depends on the size of the medium-sized objects to be preserved. As this can vary a lot from image to image, the user must keep adjusting this parameter until the output image looks good.

Moreover, we have added a parameter that increases the negative amplitudes of the band pass filtered image, and another one that increases the positive amplitudes. In this way, we can choose to what extent the white and black objects shall be preserved, respectively. These two parameters should also be gradually changed by the user until a satisfactory image is obtained. If there are no medium-sized objects, these parameters are set to zero, and should therefore be default values.

What happens if the presented method is applied to an image that has very little illumination differences, or perhaps even none? The image should, of course, remain more or less unchanged, as shown in Fig. 12 in the top row. The input image to the left is processed by the method with the previously mentioned parameters set to zero. Once again, we get the washed-out appearance of the medium-sized objects in the middle image. To the right, both these parameters are set to 1.0. A cut section through the image running through the black object is shown at the bottom, where the left (input image) is the blue curve, the middle image (corrupted) is the green curve, and the rightmost image is represented by the red curve, which is close to the blue curve. The reason why it is not identical to the blue curve is that the band pass filter cannot extract the objects perfectly. We have not tried to find the best band pass filter in this paper, but rather to show the main idea. The filter used was a simple DOG filter, and we are aware that other band pass filters might give an even better result.

It should be noted that the checkerboard test image is rather extreme in view of its square pulse content, compared to aerial photos which contain more variations. Figure 13 shows a cut section (going downwards from the top) of the rightmost image from the top of Fig. 9. The top graph, where the uneven illumination is clearly visible, represents the original image. The middle graph shows the result of applying the original method with




Fig. 12. Top: An image with perfectly even illumination is shown on the left, while the middle and rightmost images show the results obtained after processing it with the considered algorithm for different settings of the parameters affecting preservation of medium-sized objects. Bottom: The images from left to right are represented by blue, green and red curves, respectively, which show a cut through the black object. The left and right images become very similar, as indicated by the blue and red curves.

no preservation of medium-sized objects. That image is depicted in Fig. 9, on the right in the middle row. The overall illumination is improved, which can be seen by comparing the blue line with the red line (inserted for the sake of comparison). The bottom graph shows the result of applying the new method (the output image can be seen in Fig. 9, on the right in the bottom row). We can see there how the river Arno, represented by the low flat area to the right, is now much darker compared to the original method. The green line has been inserted to make the comparison easier, and the dark river now cuts through the green line, while it was lifted up by the original method. We can also note that the darker field at the top of the input image has become a bit darker after applying the proposed method, just as it should (see the leftmost part of the curve).




Fig. 13. Top: a cut section of the original image (see Fig. 9 on the right). Middle: the same cut section after applying the original method. Bottom: The proposed method has been applied, and the river (flat area to the right) is darker, as it should be, i.e. it is not lifted up above the inserted green reference line.

The difference between the middle and bottom curves might seem very small, but the difference in appearance is really noticeable in Fig. 9. We can also note that the sides of the river are not lifted up in the bottom curve as much compared to the red line as is the case in the middle curve. Hence, the annoying ripple effect is removed by the new method. This ripple effect is more visible in the image on the left in Fig. 9, where it is noticeable as a darker edge around the whiter field.

Finally, we have used one parameter to avoid contrast enhancement in already high contrast areas, as described at the end of section 3.1. Hence there are a number of parameters to be adjusted by the user. However, they can be set to standard default values, as explained, making it easy for the user to then adjust them in order to obtain an optimal image.

The proposed method is relatively fast in matlab (2.4 s for a scale factor of 10 and 3 iterations for the leftmost image in Fig. 9). However, a GPU implementation would make it possible to adjust parameters more or less in real time even for such large images as we have used here, making it suitable also for commercial applications.

6. Conclusions

The proposed algorithm uses luminance mapping for illumination correction by computing the local mean and deviation. These are obtained by Gaussian filtering, which is a computationally costly process. Therefore, we have proposed to carry out these operations on downsampled versions of the input image. The resulting blurred images are subsequently upsampled to full size with little loss of accuracy. Large objects in the images, having a different intensity than their surroundings, are removed by band pass filtering in order to avoid their washing out, which is what illumination correction generally does. They can be preserved by removing them before correction and adding afterwards. By iterating this process more than once, including removing, correction and adding, a better result is obtained as the band pass filtered image tends to contain some low variations. Several examples shown in the paper prove that the method can be used for both gray scale and colour images with good results.

References

- 1968**
- [1] Oppenheim, A.; Schafer, R.; Stockham, T., Jr.: Nonlinear filtering of multiplied and convolved signals. *Audio and Electroacoustics, IEEE Transactions on*, Volume: 16 Issue:3, pp. 437-466.
- 1983**
- [2] Burt, P., J. and Adelson, E., H.: A multiresolution spline with application to image mosaics. *Journal ACM Transactions on Graphics (TOG)*, Volume 2 Issue 4.
- [3] Fahnestock, J. D. and Schowengerdt, R.A.: Spatially variant contrast enhancement using local range modification, *Optical Engineering* 22, pp. 378-381.
- 1987**
- [4] Pizer, S. M., Amburn, E. P., Austin, J. D., Cromartie, R., Geselowitz, A., Greer, T., Romeny, B. T. H., Zimmerman, J. B. and Zuiderveld, K.: Adaptive histogram equalization and its variations. *Computer Vision, Graphics, and Image Processing*. Volume 39, Issue 3, September, pp. 355-368.
- 1994**
- [5] Puff, D. T., Pisano, E. D., Muller, K. E., Johnston, R. E., Hemminger, B. M., Burbeck, C. A., McLelland, R. and Pizer, S. M.: A Method for Determination of Optimal Image Enhancement for the Detection of Mammographic Abnormalities. *Journal of Digital Imaging*, Vol 7, No 4 (November), pp 161-171.
- [6] Vrhel, M. J. and Trussell, H. J.: Filter considerations in color correction. *IEEE Trans. Image Processing*, vol. 3, pp. 147-161, Mar.
- 1995**
- [7] Laine, A.; Jian Fan; Wuhai Yang: Wavelets for contrast enhancement of digital mammography. *Engineering in Medicine and Biology Magazine, IEEE*, Volume: 14 Issue:5. pp. 536 - 550.
- 1997**
- [8] Reeves, T.H., Jernigan, M.E.: Multiscale-based image enhancement. *Electrical and Computer Engineering. IEEE Canadian Conference on*. Issue Date: 25-28 Vol 2. pp. 500-503.
- 2000**
- [9] Young, I. T. Shading Correction: Compensation for Illumination and Sensor Inhomogeneities, *Current Protocols in Cytometry*. pp 1-12.
- 2001**
- [10] Hertzmann, A., Jacobs, C. E., Oliver, N., Curless, B, and Salesin, D. H.: Image Analogies. *SIG-GRAPH 2001 Conference Proceedings*. pp. 327-340.

2002

- [11] Ko, J., Kim, E. and Byun, H.: A simple illumination algorithm for face recognition. Proceeding of PRICAI, pp. 532-541.
- [12] Sidney F. R.: Applied photographic optics, 3rd ed., Focal Press.
- [13] Yoon, J. H., Ro, Y. M.: Enhancement of the Contrast in Mammographic Images, using the Homomorphic Filter Method. INF. & SYST., VOL.E85-D, NO.1 JANUARY LETTER. pp. 298-303.

2003

- [14] Leong F. J. W-M., Brady, M., O'D McGee J.: Correction of uneven illumination (vignetting) in digital microscopy images J Clin Pathol. pp. 619 - 621.
- [15] Pérez, P., Gangnet, M. and Blake, A.: Poisson Image Editing, Journal ACM Transactions on Graphics (TOG) - Proceedings of ACM SIGGRAPH, Volume 22 Issue 3, pp. 313-318
- [16] Tzelepis, N. and Nakos, B.: A Study on the Lighting Factors affecting Relief Presentation, Proceedings of the 21st International Cartographic Conference (ICC), pp. 1343-1350.
- [17] Zhu, J., Liu, B. and Schwartz, S. C. General illumination correction and its application to face normalization. Proceeding of ICASSP. pp 133-136.

2004

- [18] Levin, A., Zomet, A., Peleg, S., Weiss, Y.: Seamless Image Stitching in the Gradient Domain. In ECCV (4) pp. 377-389.
- [19] Yu. W. Practical anti-vignetting methods for digital cameras. IEEE Trans. on Cons. Elect., 50. Pp. 975-983.

2005

- [20] Goldman, D. B. and Chen, J. H.: Vignette and Exposure Calibration and Compensation. Proceedings of ICCV '05, pp. 89-906.
- [21] Pajares, G., Ruz, J. J. and de la Cruz, J. J.: Performance Analysis of Homomorphic Systems for Image Change Detection. Pattern Recognition and Image Analysis, LNCS, Volume 3522/2005, pp. 563-570.
- [22] Rocchini, D. and Di Rita, A.: Relief effects on aerial photos geometric correction. Applied Geography 25. pp. 159-168.
- [23] Yucong Guo, Xingming Zhang, Huangyuan Zhan, and Jing Song, A Novel Illumination Normalization Method for Face Recognition. Advances in Biometric Person Authentication, International Workshop on Biometric Recognition Systems, IWBRIS 2005, Beijing, China, Volume 3781, pp. 23-30.

2006

- [24] Zheng, Y., Lin, S., Kang. S. B. : Single-Image Vignetting Correction. Computer Vision and Pattern Recognition, IEEE Computer Society Conference. pp. 461-468.

2007

- [25] Agus, O.; Ozkan, M.; Aydin, K.: Elimination of RF Inhomogeneity Effects in Segmentation, Proceedings of the 29th Annual International Conference of the IEEE EMBS, pp. 2081-2084.
- [26] Ardizzone, E., Pirrone, R., La Bua, S., Gambino, O.: Volumetric Bias Correction. Computer Vision/Computer Graphics Collaboration Techniques: MIRAGE 2007. LNCS, vol. 4418, pp. 525-533.

2008

- [27] Ashiba, H. I., Awadalla, K. H., El-Halfawy S. M., and Abd El-Samie F. E.-S.: Homomorphic enhancement of infrared images using the additive wavelet transform. Progress In Electromagnetics Research C, Vol. 1, pp. 123-130.
- [28] Sonka, M., Hlavac, V. and Boyle, R.: Image Processing, Analysis, and Machine Vision. Thomson Learning, USA, pp.38.

2010

- [29] ERDAS Field Guide, pp. 525 (2010)
http://www.erdas.com/Libraries/Tech_Docs/ERDAS_Field_Guide.sflb.ashx

2011

- [30] Hast, A.; Marchetti, A.: Retrospective Illumination Correction of Greyscale Historical Aerial Photos. International Conference on Image Analysis and Processing, Poster. pp. 275-284.
- [31] Russ, J.C.: The Image Processing Handbook, Sixth Edition, CRC Press. pp. 370.