

Retrospective Illumination Correction of Greyscale Historical Aerial Photos

Anders Hast* and Andrea Marchetti

Consiglio Nazionale delle Ricerche (CNR), Institute of Informatics and Telematics (IIT)
Via Moruzzi, 1 – CNR, Research Area
Pisa, Italy
{Anders.Hast,Andrea.Marchetti}@iit.cnr.it

Abstract. Illumination correction is a method aiming at removing the influence of light from the environment and other distorting factors in the image capture process. A novel algorithm based on luminance mapping is proposed that both removes the low frequency variations in intensity as well as increases the contrast in low contrast areas. Moreover, it avoids the common problems with homomorphic filters. This algorithm is being applied on historical aerial photos with good results.

Keywords: Illumination Correction, Luminance Mapping, Image Stitching, Image Mosaicing, Vignetting.

1 Introduction

In many disciplines dealing with images the problem of varying illumination of the scene or object must be faced. Hence both the problem and the remedy have got different names depending on discipline. Vignetting [24, 26] occurs due to different mechanisms [7], which causes a brightness falloff away from the image centre and this is prevalent in photography. Non uniform illumination or even dirty lenses and dust [23] causes problems in Microscopy [11]. Such procedure is often called illumination or shading correction and in this paper we will refer to it as *Illumination Correction*. In Magnetic Resonance Imaging (MRI) the varying shade is known as RF-inhomogeneity or bias [1,2]. In Face Recognition the varying illumination is a challenging problem [25, 27] and the 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 [17].

In aerial and satellite photos the physical lighting also affects the relief presentation [19, 20]. 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 [4, 12, 15].

* This work was carried out during the tenure of an ERCIM "Alain Bensoussan" Fellowship Programme.

In this paper we will focus on a method that we developed for *retrospective illumination correction of greyscale historical aerial photos* that will not only minimize the variance in illumination over the image but can also be used to increase contrast in low contrast areas in the image.

1.1 Main Methods

As mentioned in the introduction the same problem of uneven illumination occurs in many fields and the referenced papers suggest different methods specialized for the application in question. However there are some fundamental aspects that will be shortly discussed here.

Theoretically, the intensity in the image can be divided into two components, the illumination of the object and the amount of light reflected by the object [14]. The homomorphic filter [13] aims at removing the illumination, characterized by low spatial variations, from the image and keeping reflectance, characterized by high frequency changes. In terms of the *illumination* $L(x,y)$ and *reflectance* $R(x,y)$, the *intensity* $H(x,y)$ of a pixel can be modelled as:

$$H(x, y) = L(x, y)R(x, y). \quad (1)$$

The homomorphic filter makes use of the Fourier transform to remove the *intensity* and hence it is necessary to convert the multiplication to addition, which can be done by taking the log of the functions, as the Fourier transform only can be used when the noise is an additive term. This approach has some disadvantages, for instance the necessary image size padding when the image size is not a power of 2, which can lead to distortions in the boundary regions [5]. Another problem with this approach is that it will merely increase contrast in low contrast regions while the main low frequencies remains. A better way to look at the problem [27] is to model the lighting change as a local affine transformation (AT) of the pixel value:

$$H(x, y) = A(x, y)h(x, y) + B(x, y). \quad (2)$$

where $h(x,y)$ is the original image and $H(x,y)$ is the captured image as a result of the local lighting, which thus have both and multiplicative and additive effect on the result. Zhu et al [27] removes the low spatial variations caused by $B(x,y)$ using a multi resolution low pass filter that estimates the background lighting. Leung et al [11] takes the same approach using a Gaussian filter:

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

where σ defines the width of the distribution used to delimit the low frequencies. The larger the value of σ , the greater the smoothing effect. Others, for an example Yu [24] handle the background light using a reference picture acquired in a controlled environment with known lighting conditions and a white object. This is of course not always possible, especially not in retrospect.

The gain $A(x,y)$ in eq. 2 is handled by a homomorphic filter and this will increase the contrast. The approach by Zhu gives better results but the main disadvantages with the homomorphic filtering remains. The resulting image computed by AT is:

$$\check{h}(x, y) = \frac{\kappa}{\check{A}(x, y)} \left(H(x, y) - \check{B}(x, y) \right). \quad (4)$$

Where $\check{A}(x, y)$, $\check{B}(x, y)$ are the estimate of $A(x, y)$, $B(x, y)$ respectively and κ is a constant. Zhu et al modifies this equation further but the equation shown here is the basis. Yucong et al [25] have improved the AT by computing the estimated gain as a low pass filtered estimate:

$$\check{A}(x, y) = \exp \left(\text{LPF} \left(\ln \left(h(x, y) - \check{B}(x, y) \right) \right) \right). \quad (5)$$

where *LPF* is a low pass filter. They elaborate further on his approach using histogram equalization and also combining with ICR [9] explained shortly in next subsection, however these are the basic ideas.

1.2 Other Approaches

Others have developed methods for finding the *vignetting function* from a single image [26] or from multiple images, with or without the response curve [7]. The ICR (Illumination Compensation based on Multiple Regression Model) [9] is aiming to find the plane that best fits the intensity distribution of the image using the multiple regression model. Then this plane is used to remove the illumination of face images. The Local Range Modification (LRM) [6], finds the interpolated minimum and maximum pixel values in a neighbourhood (contextual region) and stretches them to the desired range via the equation:

$$Y(x, y) = \frac{C}{\max - \min} (X(x, y) - \min). \quad (6)$$

where C is a constant. Similarly the contrast-limited adaptive histogram equalization method (CLAHE) [16], stretches the histogram in its contextual region. This improves the image in such way that the contrast is enhanced. Young et al [25] make use of homomorphic filters and Morphological filtering in order to correct images in microscopy. Here the morphological filter is used to obtain an estimate of the background illumination.

Wavelet based approaches for image and contrast enhancement are also popular [18]. Laine et al [10] use wavelets to enhance contrast in digital mammography. Yet others make use of a combination of methods as wavelets and the homomorphic filter [3,22].

1.3 Luminance Mapping

Colour correction can be used to both correct and filter colours in an image [21]. One important method is Luminance mapping [8]. The idea is to apply a linear map that matches the means and variances of the 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 :

$$A(x, y) = \frac{\sigma_B}{\sigma_A} (A(x, y) - \mu_A) + \mu_B. \quad (7)$$

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, if applied an all three colour channels, is that image A will have similar looking colours as image B . The result is shown in Fig. 1 where the original Lena and Mandrill images are to the left and the luminance mapped images to the right where Lena has got the colours from the Mandrill and vice versa. This method has not got any obvious connection to illumination correction, however it is the basis for the proposed method discussed in the next section.



Fig. 1. The original Lena and Mandrill images to the left and the luminance mapped images to the right, where Lena has got the colours from the Mandrill and vice versa.

2 The Proposed Method

In the approach by Zhu et al the low frequencies are removed first by subtracting $\check{B}(x, y)$, which is an estimation of the background lighting and then the contrast is increased by multiplication of $\kappa/\check{A}(x, y)$ (Note that they use a more elaborated version of this factor in their paper, including eq. 5). This implies that $\check{B}(x, y)$ contains the low frequencies in the image. If eq. 7 is applied using, not an overall estimate of the mean but an estimate of the local mean μ it is possible to obtain a similar result. Let us take the image of a checkerboard in the left of Fig. 2 as an example, which has an obvious intensity variation. Looking at a horizontal cut section taken from the centre, depicted to the right, we can see the intensity variation from left to right. The *local mean* is the red curve that runs in the centre of the curve. If the *local mean* is translated to the *global mean*, i.e. pushing down the curve so that it becomes a line, then a big portion of the varying light is removed.

A *low pass filter* is used to remove an estimate of the illumination, which are the lowest frequencies in the image and thus an approximation of the *local mean* $\mu_A(x, y)$. This is done by using a Gaussian filter (eq. 3) modified in a way such as it also compensates for the borders. However the mask must be quite large, up to a size of the input image. Even with modern computers it becomes impractical for images of sizes up to 5000x6000 pixels, which we are dealing with in the case of historical aerial photos. This problem can be easily overcome by the fact that the estimate is a heavily blurred version of the input image and therefore we can compute $\check{\mu}_A(x, y)$ using a downsampled and blurred version $\check{H}(x, y)$ of the image and then upscale it to $\check{\mu}_A(x, y)$ using a bilinear interpolation. The resulting blurred image will be very close to what the full scale Gaussian filter would produce and will not be detrimental for the final quality, however in Fig. 2-4 a scale factor of 1 was used, i.e. no downsampling at all.




Fig. 2. The image to the left has an obvious shading artefact visible as an intensity fall off away from the centre. The diagram to the right shows a cut section in the middle in blue and the local mean in red.

Fig. 3 shows how the image in Fig.2 have been processed by the simple method of subtracting the local mean estimated by the blurred image and adding the global mean, in order to correct the illumination. It can be seen that the curve is corrected, however the contrast is low on the sides. In this case the image has been processed three times with varying size of the Gaussian mask (in sizes of the whole image: $\frac{1}{4}$, $\frac{1}{2}$ and 1).




Fig. 3. After processing the illumination is corrected as also is shown to the left in the cut section. However, the contrast is not yet corrected.

In the next step the local deviation is computed as the difference between the local mean taken from the blurred image and the pixel values in a block region, whose size is equal to the Gaussian kernel. The block size was chosen the same as for the computation of the local mean for practical reasons, however there is nothing that prevents for testing other sizes. The local deviation $\tilde{\sigma}_A(x, y)$ is hence computed as the absolute value of the difference between $\tilde{\mu}_A(x, y)$ and the downscaled version $\tilde{H}(x, y)$ of the original image (x, y) :

$$\tilde{\sigma}_A(x, y) = \left| \tilde{H}(x, y) - \tilde{\mu}_A(x, y) \right| \tag{8}$$

The result $\tilde{\sigma}_A(x, y)$ is then subsequently blurred using the Gaussian mask and then upsampled to its full size $\check{\sigma}_A(x, y)$. Therefore this algorithm is doing something similar as Yucong et al as it appears in eq. 5, however, with different mask sizes and without

the logarithm and exponential function. Moreover it is used in a different and more simple way as it is plugged in to eq. 7 and we get:

$$\check{h}(x, y) = \frac{\sigma_B}{\check{\sigma}_A(x, y)} (H(x, y) - \check{\mu}_A(x, y)) + \mu_B. \tag{9}$$

This equation is used in three steps as before with the same sizes of the Gaussian kernel and the result is shown in Fig. 4. The proposed contrast enhancement using luminance mapping yields a straightened curve where the amplitude is almost the same over the whole curve and the shading artefact is removed from the image.





Fig. 4. After further processing the contrast is enhanced and the two halves become almost identical

3 Results

Let us examine some results of using the proposed method and look at some necessary changes to the algorithm. A historical photo showing the river Arno close to Pisa (Italy) taken during World War II is shown in Fig 5. On the top left is a cut out from the original photo (3800x3400 pixels). The illumination is obviously unevenly spread over the picture. On the top right is the process of pushing down the mean done three times, in the same manner as for the checkerboard in Fig.2-4. Since the light was either very low in the top or very high in the bottom, these regions will have quite low contrast. The change in the local mean of the image does not affect the contrast and therefore the contrast enhancement has been applied in the bottom row. The contrast is now higher in the top and in the bottom right of the image. However, in the same time as low contrast regions are getting increased contrast, the algorithm will also decrease the already high contrast areas. Hence it is often necessary to avoid contrast enhancement in this areas. This can be easily done by not allowing $\check{\sigma}_A$ to be larger than σ_B , or not larger than some percentage of σ_B . In this manner we can adjust the image to obtain the desired result. The bottom left shows the result when high contrast regions are kept without lowering the contrast. A scale factor of 10 was used processing these images.

Moreover, the reason for repeating the process three times, starting with a Gaussian mask of size $\frac{1}{4}$ of the image and then doubling the size, is evident in this picture because, if we would have taken a smaller size, let us say $\frac{1}{8}$ of the image then the river Arno would become much brighter as the local mean would be closer to the true



© MiBAC-ICCD, Aerofototeca Nazionale, fondo RAF.

Fig. 5. A historical photo showing the river Arno close to Pisa taken during WW- II. On the top left is a cut out from the original photo (3400x3800 pixels), with unevenly spread illumination. On the top right is the mean corrected three times. Low contrast regions are apparent in the top and bottom of the image where the light was low or high. On the bottom row left, is the contrast increased, which also means that high contrast regions will have lower contrast. The bottom right shows the result when high contrast regions are kept without lowering the contrast.

value. This can be seen in Fig. 6, where the mask size is $\frac{1}{4}$ to the left and $\frac{1}{8}$ to the right. This is a vertical cross section of the top left image in Fig. 5. Clearly the red curve in the right picture is a better approximation of the blue curve, however it will lift up Arno, which is the dip in the curve to the right. Hence, it is necessary to use a larger mask size resulting in an image that is not fully corrected. Therefore, it is often necessary to repeat the process with a larger mask for each step to correct this behaviour. Experimentally we found that doubling the mask is generally a good approach for obtaining visually pleasing results.





Fig. 6. A vertical cross section of the top left image in Fig. 5, where the mask size is $\frac{1}{4}$ to the left and $\frac{1}{8}$ to the right. A smaller mask gives a local mean (in red), which is a better approximation of the curve (in blue). However, it will also reduce the contrast in the image.

In fig. 7. There is another historical image taken during WW-II close to Pisa, showing an open field in the bottom with a forest surrounding it. The illumination is once again very uneven and the corrected result to the left has an even illumination.



© MiBAC-ICCD, Acrofitoteca Nazionale, fondo RAF.

Fig. 7. A field close to Pisa during WW-II. To the left a cut out of the original image (3300x3800 pixels) and the corrected image to the right.

One can note a couple of things, for instance that the cut was done wrong in the sense that there is some parts of an arrow visible in the top that belongs to the border of the original image. Without the illumination correction it would have been hard for the human eye to see this. Moreover, the field becomes quite dark as the algorithm tend to set the even light over the whole picture and the field occupies a substantial part of it. Setting a smaller Gaussian mask helps a little but then the light will be less even in the image.

4 Conclusions

The proposed use of luminance mapping for illumination correction has not to our knowledge been done before. The other main contributions in this paper is that we propose a modification where both the local mean and deviation is obtained by a Gaussian filter, where both can be computed from downscaled versions of the original image and then upscaled by bilinear interpolation. This will decrease computation time as the Gaussian interpolations in computationally expensive, however it will not affect the result for moderate downscales. Moreover, it was shown that the contrast can easily be enhanced only where needed. The example images shows that the proposed approach of using a modified version of luminance mapping for illumination correction works very well for retrospective illumination correction of historical aerial photos.

References

1. Agus, O., Ozkan, M., Aydin, K.: Elimination of RF Inhomogeneity Effects in Segmentation. In: Proceedings of the 29th Annual International Conference of the IEEE EMBS, pp. 2081–2084 (2007)
2. Ardizzone, E., Pirrone, R., La Bua, S., Gambino, O.: Volumetric Bias Correction. In: Gagalowicz, A., Philips, W. (eds.) MIRAGE 2007. LNCS, vol. 4418, pp. 525–533. Springer, Heidelberg (2007)
3. Ashiba, H.I., Awadalla, K.H., El-Halfawy, S.M., Abd El-Samie, F.E.-S.: Homomorphic enhancement of infrared images using the additive wavelet transform. Progress In Electromagnetics Research C 1, 123–130 (2008)
4. Burt, P.J., Adelson, E.H.: A multiresolution spline with application to image mosaics. *Journal ACM Transactions on Graphics (TOG)* 2(4) (1983)
5. ERDAS Field Guide, p. 525 (2010), http://www.erdas.com/Libraries/Tech_Docs/ERDAS_Field_Guide.sflb.ashx
6. Fahnestock, J.D., Schowengerdt, R.A.: Spatially variant contrast enhancement using local range modification. *Optical Engineering* 22, 378–381 (1983)
7. Goldman, D.B., Chen, J.H.: Vignette and Exposure Calibration and Compensation. In: Proceedings of ICCV 2005, pp. 89–906 (2005)
8. Hertzmann, A., Jacobs, C.E., Oliver, N., Curless, B., Salesin, D.H.: Image Analogies. In: SIGGRAPH 2001 Conference Proceedings, pp. 327–340 (2001)
9. Ko, J., Kim, E.-J., Byun, H.: A simple illumination normalization algorithm for face recognition. In: Ishizuka, M., Sattar, A. (eds.) PRICAI 2002. LNCS (LNAI), vol. 2417, pp. 532–541. Springer, Heidelberg (2002)
10. Laine, A., Fan, J., Yang, W.: Wavelets for contrast enhancement of digital mammography. *IEEE Engineering in Medicine and Biology Magazine* 14(5), 536–550 (1995)
11. Leong, F.J.W.-M., Brady, M., O'D McGee, J.: Correction of uneven illumination (vignetting) in digital microscopy images. *J. Clin. Pathol.*, 619–621 (2003)
12. Levin, A., Zomet, A., Peleg, S., Weiss, Y.: Seamless image stitching in the gradient domain. In: Pajdla, T., Matas, J(G.) (eds.) ECCV 2004. LNCS, vol. 3024, pp. 377–389. Springer, Heidelberg (2004)
13. Oppenheim, A., Schaffer, R., Stockham Jr., T.: Nonlinear filtering of multiplied and convolved signals. *IEEE Transactions on Audio and Electroacoustics* 16(3), 437–466 (1968)

14. Pajares, G., Ruz, J.J., de la Cruz, J.M.: Performance analysis of homomorphic systems for image change detection. In: Marques, J.S., Pérez de la Blanca, N., Pina, P. (eds.) *IbPRIA 2005*. LNCS, vol. 3522, pp. 563–570. Springer, Heidelberg (2005)
15. Pérez, P., Gangnet, M., Blake, A.: Poisson Image Editing, *Journal ACM Transactions on Graphics (TOG)*. In: *Proceedings of ACM SIGGRAPH*, vol. 22(3), pp. 313–318 (2003)
16. Pizer, S.M., Amburn, E.P., Austin, J.D., Cromartie, R., Geselowitz, A., Greer, T., Romeny, B.T.H., Zimmerman, J.B., Zuiderveld, K.: Adaptive histogram equalization and its variations. In: *Computer Vision, Graphics, and Image Processing*, vol. 39(3), pp. 355–368 (September 3, 1987)
17. Puff, D.T., Pisano, E.D., Muller, K.E., Johnston, R.E., Hemminger, B.M., Burbeck, C.A., McLelland, R., Pizer, S.M.: A Method for Determination of Optimal Image Enhancement for the Detection of Mammographic Abnormalities. *Journal of Digital Imaging* 7(4), 161–171 (1994)
18. Reeves, T.H., Jernigan, M.E.: Multiscale-based image enhancement. In: *IEEE Canadian Conference on Electrical and Computer Engineering*, Issue Date: 25-28, vol. 2, pp. 500–503 (1997)
19. Rocchini, D., Di Rita, A.: Relief effects on aerial photos geometric correction. *Applied Geography* 25, 159–168 (2005)
20. Tzelepis, N., Nakos, B.: A Study on the Lighting Factors affecting Relief Presentation. In: *Proceedings of the 21st International Cartographic Conference (ICC)*, pp. 1343–1350 (2003)
21. Vrhel, M.J., Trussell, H.J.: Filter considerations in color correction. *IEEE Trans. Image Processing* 3, 147–161 (1994)
22. Yoon, J.H., Ro, Y.M.: Enhancement of the Contrast in Mammographic Images, using the Homomorphic Filter Method. *Inf. & Syst. Letter E85–D(1)*, 298–303 (2002)
23. Young, I.T.: Shading Correction: Compensation for Illumination and Sensor Inhomogeneities. *Current Protocols in Cytometry*, 1–12 (2000)
24. Yu, W.: Practical anti-vignetting methods for digital cameras. *IEEE Trans. on Cons. Elect.* 50, 975–983 (2004)
25. Guo, Y., Zhang, X., Zhan, H., Song, J.: A novel illumination normalization method for face recognition. In: Li, S.Z., Sun, Z., Tan, T., Pankanti, S., Chollet, G., Zhang, D. (eds.) *IWBRIS 2005*. LNCS, vol. 3781, pp. 23–30. Springer, Heidelberg (2005)
26. Zheng, Y., Lin, S., Kang, S.B.: Single-Image Vignetting Correction. In: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 461–468 (2006)
27. Zhu, J., Liu, B., Schwartz, S.C.: General illumination correction and its application to face normalization. In: *Proceeding of ICASSP*, pp. 133–136 (2003)