SINGLE IMAGE VIGNETTING CORRECTION WITH NATURAL IMAGE STATISTICS IN DERIVATIVE DOMAINS

Siwei Lyu

Computer Science Department
University at Albany, SUNY
Albany, NY 12222, USA

ABSTRACT

Vignetting is the phenomenon of reduced brightness in an image at the peripheral region compared to the central region. In this paper, we describe a new method for model based single image vignetting correction. We use the statistical properties of natural images in the discrete derivative domains and formulate the vignetting correction problem as a maximum likelihood estimation. We further provide a simple and efficient procedure for better initialization of the numerical optimization. Empirical evaluations of the proposed method using synthesized and real vignetted images show significant gain in both performance and running efficiency.

1. INTRODUCTION

Due to the design of camera lenses, volume of light transmitted to the image sensor tends to decrease somewhat from the center of the image to the corners, and such a phenomenon is known as *vignetting* [8]. Vignetting is also intentionally added, by burning the outer edges of the photograph for film stock or using digital imaging techniques such as the *Lens Correction* filter in Photoshop, to achieve certain creative effects or to draw attention to the main subject. However, there are many situations in photography such as wide-angle landscapes where vignetting is unwanted. The attenuation of brightness in a vignetted image is also undesirable for computer vision algorithms where accurate intensity values are needed. On the other hand, the specific way a camera lens causes vignetting can provide useful ballistic information for its identification, and is useful for forensic analysis [1].

In this paper, we describe a new method for model based single image vignetting correction, which aims to factorize a vignetted image into the product of the original image and a spatially varying vignetting function, using only one observed vignetted image and a parametric model of the vignetting function. This is an ill-posed problem, as many different combinations of the original image and the vignetting function can produce the same observed vignetted image. To obtain a reasonable solution, we use statistical properties of natural photographic images that do not have vignetting effects. Specifically, using marginal distributions of the *mean* responses of natural images in the derivative domains, we formulate the vignetting correction problem as a maximum likelihood estimation. Compared to the existing works on single image vignetting correction [13, 14], the proposed work has the following important properties: (i) our method does not require the vignetting function to be centered at the image center – it estimates the center from the vignetted image as part of the procedure. (ii) Our method handles elliptical vignetting functions. (iii) We use maximum likelihood estimation based on a robust statistical model for natural images, which leads to a simpler optimization procedure. Empirical evaluations of the proposed method using both synthesized and real vignetted images show significant gain in both performance and running efficiency.

2. RELATED WORK

Correcting vignetting is easy if the vignetting function is known. For high-end single lens reflex digital cameras, such information can be stored on chip for stock lenses and used to correct vignetting after an image is captured, an example is the *Peripheral Illumination Correction* function in the recent Cannon EOS 50D cameras. If the vignetting parameters are not directly accessible, they can still be obtained by camera calibration with an image of uniform brightness, e.g., [5, 12]. If the camera is not at our disposal, the vignetting function can be estimated using a sequence of images of a static scene captured by the same camera and lens [3, 4, 6].

The most challenging case for vignetting correction, however, is when we only have a single vignetted image and no specific information about the camera and lens. A single image vignetting correction method was described in [13], where the vignetted image is first segmented into texture and non-texture regions, and the vignetting function is estimated from the non-texture regions with a robust outlier exclusion. However, the requirement of image segmentation makes this algorithm slow and segmentation errors affect the final correction result. An improved method was described in [14] based on the empirically observed symmetry in the distributions of directional derivatives. As no segmentation is required, this method achieves improvement in both performance and running time. On the other hand, both methods assume that (1) the center of the vignetting function is fixed at the image center, and (2) the shape of the vignetting function is circular. Thus they are not able to deal with vignetted images cropped off-center or from elliptical vignetting functions.

3. SINGLE IMAGE VIGNETTING CORRECTION

Let us introduce some notations and the basic setting of the vignetting correction problem. First, we assume that a vignetted image $I_v(x, y)$ is the product of the original image $I(x, y)$ and the vignetting function $\phi(\mathbf{v}; \theta)$, $I_v(\mathbf{v}) = I(\mathbf{v})\phi(\mathbf{v}; \theta)$, where $\mathbf{v} = (x, y)^T$ is the vector of pixel locations, with the origin of the image coordinate system being the image center, and θ concisely denote all parameters of ϕ. Taking logarithm on both sides, this is equivalent to

$$i_v(\mathbf{v}) = i(\mathbf{v}) + \log A(\mathbf{v}) + \log G(\mathbf{v}), \tag{1}$$

with $i(\mathbf{v})$ and $i_v(\mathbf{v})$ standing for $\log I(\mathbf{v})$ and $\log I_v(\mathbf{v})$, respectively. The goal of model based single image vignetting correction is to obtain parameters in the vignetting function, θ, and recover the original image $I(\mathbf{v})$, using only a single observed vignetted image $I_v(\mathbf{v})$. For the vignetting function, we adopt a generalization of the Kang-Weiss vignetting model [5], which allows for elliptical shapes for the vignetting functions, as

$$\phi(\mathbf{v}; \theta) = A(\mathbf{v}; \mathbf{c}, P)G(\mathbf{v}; \{\alpha_i\}_{i=1}^k). \tag{2}$$

For conciseness, we will use $A(\mathbf{v})$ and $G(\mathbf{v})$ instead to denote these functions subsequently. $A(\mathbf{v})$ is the illumination factor and is defined as $A(\mathbf{v}) = 1/(1 + r^2/2)^2$, where

$$r = \sqrt{(\mathbf{v} - \mathbf{c})^T P(\mathbf{v} - \mathbf{c})}, \mathbf{c} = \begin{pmatrix} c_1 \\ c_2 \end{pmatrix}, P = \begin{pmatrix} p_1 & p_2 \\ p_2 & p_3 \end{pmatrix},$$

are the Mahalanobis distance to the center, the center and the shape matrix of the vignetting function, respectively. The effective focal length f of the vignetting function is implicit in the shape matrix P, and for a circular vignetting function, one has $p_1 = p_3 = 1/f^2$ and $p_2 = 0$. $G(\mathbf{v})$ is the geometric factor of the vignetting function and we use a polynomial of r as $1 - \sum_{i=1}^k \alpha_i r^i$, for $k = 5$ as in [14].

We solve the vignetting correction problem by transforming Eq.(1) into the discrete derivative domains. We use D^r_Δ to denote discrete derivative operators used in this work, where $t \in \{x, y, xx, xy, yy\}$ signifies the type of derivative, and Δ gives the step size in computing such derivatives. Specifically, we use first order derivative operators: $D^x_\Delta I(x, y) = I(x + \Delta, y) - I(x, y)$, $D^y_\Delta I(x, y) = I(x, y + \Delta) - I(x, y)$, and second order derivative operators:
$D^{xx}_\Delta I(x, y) = I(x + \Delta, y) + I(x - \Delta, y) - 2I(x, y)$, $D^{yy}_\Delta I(x, y) = I(x, y + \Delta) + I(x, y - \Delta) - 2I(x, y)$, and

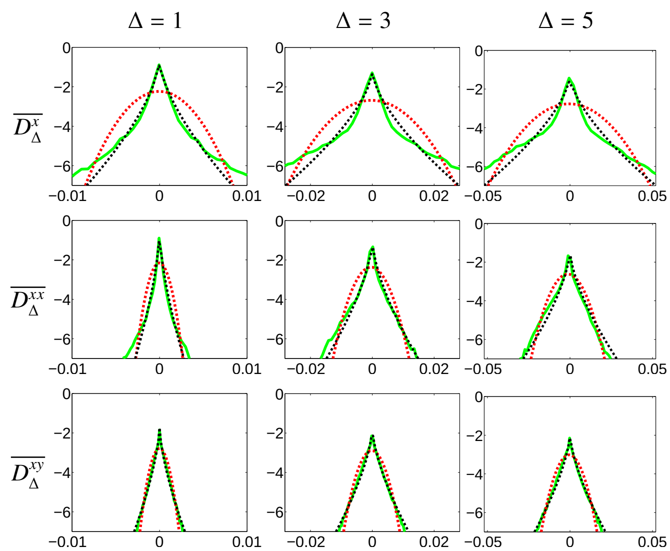

Fig. 1: *Log marginal histograms (green solid curves) for different $\overline{D^t_\Delta i(\mathbf{v})}$ and Δ obtained with randomly chosen blocks from the van Hateren database [11]. Red and black dashed curves correspond to Gaussian densities of the same mean and variance, and the optimally fitted generalized Laplacian densities, respectively.*

$D^x_\Delta I(x, y) = I(x + \Delta, y + \Delta) - I(x, y + \Delta) - I(x + \Delta, y) - I(x, y)$. When applied to an image, these discrete derivative operators can be efficiently implemented as convolution of the image with appropriate filters[1]. Conceptually, the discrete differential operators, after normalizing with the step size, implement the corresponding continuous differential operators on the discrete pixel lattice. For instance $\frac{1}{\Delta} D^x_\Delta$ corresponds to $\frac{\partial}{\partial x}$, and $\frac{1}{\Delta^2} D^{xx}_\Delta$ corresponds to $\frac{\partial^2}{\partial x^2}$.

Applying the discrete derivative operators to both sides of Eq.(1), we have, $D^r_\Delta i_v(\mathbf{v}) = D^r_\Delta i(\mathbf{v}) + D^r_\Delta \log A(\mathbf{v}) + D^r_\Delta \log G(\mathbf{v})$. We then average both sides over all pixel locations $\mathbf{v}$ to obtain

$$\overline{D^r_\Delta i_v(\mathbf{v})} = \overline{D^r_\Delta i(\mathbf{v})} + \overline{D^r_\Delta \log A(\mathbf{v})} + \overline{D^r_\Delta \log G(\mathbf{v})}. \tag{3}$$

Note that $\overline{D^r_\Delta i_v(\mathbf{v})}$ can be obtained from the observed vignetted image, $\overline{D^r_\Delta \log A(\mathbf{v})}$ and $\overline{D^r_\Delta \log G(\mathbf{v})}$ can be computed from the generalized Kang-Weiss vignetting function for a given parameter setting θ. Eq.(3) suggests that from these terms, we are able to compute $\overline{D^r_\Delta i(\mathbf{v})} = \overline{D^r_\Delta i_v(\mathbf{v})} - \overline{D^r_\Delta \log A(\mathbf{v})} - \overline{D^r_\Delta \log G(\mathbf{v})}$. In general, term $-\overline{D^r_\Delta \log A(\mathbf{v})} - \overline{D^r_\Delta \log G(\mathbf{v})}$ introduces statistical deviations to $\overline{D^r_\Delta i_v(\mathbf{v})}$ from $\overline{D^r_\Delta i(\mathbf{v})}$ unless the correct parameters for the vignetting function are chosen. As $\overline{D^r_\Delta i(\mathbf{v})}$ reflects properties of the original image without vignetting, we can seek vignetting function parameters θ that make $\overline{D^r_\Delta i_v(\mathbf{v})} - \overline{D^r_\Delta \log A(\mathbf{v})} - \overline{D^r_\Delta \log G(\mathbf{v})}$ consistent with $\overline{D^r_\Delta i(\mathbf{v})}$ for natural images.

[1] Though more elaborated filters such as those in [2] can be employed, we found the results are quite independent of the choice of filters.

3.1. Natural Image Statistics in Derivative Domains

Images resembling scenes from the physical world are loosely tagged as *natural images*, and are known to occupy only a small fraction of all possible images [9]. To solve the vignetting correction problem in Eq.(3), we need to investigate the statistical properties of the *average* derivative domain responses of natural images $\overline{D^t_\Delta i(\mathbf{v})}$. We collect 20,000 randomly chosen pixel blocks of size 200×200 from images in the Van Hateren calibrated image database [11]. These grayscale images are obtained from calibrated cameras with no obvious vignetting effects and linearized intensities. We first log transform the pixel intensity, then apply the discrete derivative operators to the log pixel intensities of each image block. We then collect the mean discrete derivative domain response for each of the 20,000 blocks and obtain their histograms. These histograms are shown in the log domain in Figure 1 (green solid curves) for different $\overline{D^t_\Delta i(\mathbf{v})}$ and Δ values (columns). For the sake of comparison, Gaussian densities with the same means and variances are shown as red dashed curves. Also shown as black dashed curves in each plot are the optimal fitting (using the method of moment matching [10]) of the generalized Laplacian density to these histograms. From this experiment, we note that the histograms of the mean responses of natural images in the discrete derivative domain, similar to those of the raw derivative domain responses, have means close to zero and show strong non-Gaussian characteristics. Furthermore, they can also be well approximated with generalize Laplacian models.

3.2. Vignetting Correction as Maximum Likelihood

Due to the vignetting effect, the statistical properties of the vignetted image in the derivative domains are different from the original image, especially, the vignetting function introduces a bias in Eq.(3). We model the marginal densities of the mean responses of original natural images in the derivative domains with generalized Laplacians, $p_{r,\Delta}(x) \propto \exp\left(-\left|x/\gamma_{r,\Delta}\right|^{\beta_{r,\Delta}}\right)$, with parameters $\beta_{r,\Delta}$ and $\gamma_{r,\Delta}$ estimated from the Van Hateren imagebase as described in Section 3.1 for each r and Δ. We then reformulate the vignetting correction as estimating parameter θ to maximize the joint (log) likelihood of $\overline{D^r_\Delta i(\mathbf{v})}$ with regards to different r and Δ. In detail, dropping irrelevant constants, we find parameter θ that minimizes:

$$\sum_{r,\Delta}\left(\frac{\overline{D^r_\Delta i_v(\mathbf{v})} - \overline{D^r_\Delta \log A(\mathbf{v})} - \overline{D^r_\Delta \log G(\mathbf{v})}}{\gamma_{r,\Delta}}\right)^{\beta_{r,\Delta}}, \quad (4)$$

with a constraint to ensure the shape matrix P to be positive definite, i.e., $p_1 p_3 > p_2^2$.

Compared to the objective function used in [14] that is based on the symmetry of directional derivatives, Eq.(4) is simpler and justified with a more comprehensive description of image statistics in the form of a statistical model. Optimization of Eq.(4) is implemented with coordinate descent, by alternating between steps that optimizes $\{\alpha_i\}_{i=1}^k$ with $\mathbf{c}$ and P fixed and that optimizes $\mathbf{c}$ and P with $\{\alpha_i\}_{i=1}^k$ fixed. Such a procedure guarantees to converge to a local minimum of Eq.(4), and each optimization step is implemented with gradient based numerical methods.

Fig. 2: *Examples of restoring synthetic vignetted images.* **Left***: vignetted image.* **Right***: restored image.*

4. EXPERIMENTS

We evaluate the performance and running efficiency of the proposed model based single image vignetting correction method. In the first set of experiments, we select 20 RGB JPEG images of 320×480 pixels from the Berkeley image segmentation database [7] that do not have obvious vignetting effects. We apply synthesized vignetting function to these images, and then use the method described in this work to recover the vignetting functions and the original image. As in this case, we have ground truth for both the vignetting functions and the original images, these experiments give us a chance to quantitatively evaluate the performance and running efficiency of the proposed method. The synthesized vignetting functions are created using the generalized Kang-Weiss model, Eq.(1) with several settings for the centers and shape matrices. We use the same set of $(\alpha_1, \cdots, \alpha_5)$ for the geometric component $G(\mathbf{v})$.

We apply these synthetic vignetting functions to the 20 test images, where the same vignetting function is applied to each RGB color channel individually. We evaluate vignetting correction performance using signal-to-noise ratio (SNR) with unit in dB (the higher the better), and use relative errors to quantify the estimated vignetting function parameters. Specifically, for P we use $\|P^\star - P\|_F/\|P\|_F$, where $\|\cdot\|_F$ is the Frobenius norm, and for $\mathbf{c}$ and $\{\alpha_i\}_{i=1}^k$, we use the ratio of l_2 error of the estimation and the l_2 norm of the actual parameter values. We report the performance and running efficiency corresponding to the three vignetting functions and averaged

SNR (dB)	error P (%)	error $\mathbf{c}$ (%)	error α (%)	time (sec)
24.4	8.9	16.3	7.8	76.4

Table 1: *Performance and running efficiency of the proposed algorithm. See text for details.*

Fig. 3: *Examples from the Berkeley database with strong vignetting effects (left) and their corrections with our method (right).*

over all 20 test images in Table 1. These results are based on an unoptimized implementation of the algorithm with MATLAB on a MacBook Pro laptop with due core Intel processor of 2.6GHz and 2G memory. Figure 2 shows several examples of the restored image. As these show, the vignetting effect has been largely removed from the restored images.

In the second set of experiments, we apply our method to some physically vignetted images. We choose several images from the Berkeley database with strong vignetting effects, and the correction results are shown in Fig.3. As these results show, in most of the cases our method can successfully remove the vignetting effect, and the result is comparable with that uses information of the source camera and lens. We observed that its ability to estimate the vignetting center is particularly important for performance for test images that have been cropped off-center. On the other hand, restored images with our method tend to be over-exposed closed to the image boundary. This is especially severe for images with complex texture regions (e.g.,middle panel of Fig.3).

5. CONCLUSION

We have presented a new model based single image vignetting correction method using the regular statistical properties of natural images. We start by log transforming the vignetted image and then working in the discrete derivative domains. Based on properties of the mean discrete derivative responses of natural images with no vignetting effect, we reformulate vignetting correction as maximum likelihood estimation of the vignetting function parameters. Compared to previous single image vignetting correction methods [13, 14], our method does not require image segmentation or the vignetting function to be centered at the image center, and it also works when the vignetting function has an elliptical shape. Empirical evaluations of the proposed method using synthesized and real vignetted images show significant gain in both performance and running efficiency. On the other hand, our method uses parametric model for the vignetting function and we are working to extend this methodology to nonparametric estimation. Also, as the vignetting function are specific for different lenses, we are currently investigating if we can use the estimated vignetting functions to identify stock lenses.

Acknowledgement

This work was supported by an NSF CAREER Award (IIS09-53373) and the University at Albany Faculty Research Awards Program (FRAP) - Category A.

6. REFERENCES

[1] H. Farid. Photo fakery and forensics. In *Advances in Computers*. 2009. (to appear).

[2] H. Farid and E. Simoncelli. Differentiation of discrete multidimensional signals. *IEEE Trans. Image Proc.*, 13(4):496–508, 2004.

[3] D. B. Goldman and J. H. Chen. Vignette and exposure calibration and compensation. In *ICCV*, 2005.

[4] J. Jia and C.-K. Tang. Tensor voting for image correction by global and local intensity alignment. *IEEE Trans. PAMI*, 27(1):36–50, 2005.

[5] S. Kang and R. Weiss. Can we calibrate a camera using an image of a flat textureless lambertian surface? In *ECCV*, 2000.

[6] A. Litvinov and Y. Y. Schechner. Addressing radiometric nonidealities: A unified framework. In *CVPR*, 2005.

[7] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *ICCV*, 2001.

[8] S. F. Ray. *Applied photographic optics*. Focal Press, 3rd edition, 2002.

[9] D. Ruderman. The statistics of natural images. *Network : Comp. in Neural Sys.*, 5:598–605, 1994.

[10] E. P. Simoncelli and E. H. Adelson. Noise removal via Bayesian wavelet coring. In *ICIP*, 1996.

[11] A. van der Schaaf and J. H. van Hateren. Modelling the power spectra of natural images: Statistics and information. *Vision Research*, 28(17):2759–2770, 1996.

[12] W. Yu. Practical anti-vignetting methods for digital cameras. *IEEE Trans. on Cons. Elect*, 50(2):975–983, 2004.

[13] Y. Zheng, S. Lin, and S. B. Kang. Single-image vignetting correction. In *CVPR*, 2006.

[14] Y. Zheng, J. Yu, S. Lin, S. B. Kang, and C. Kambhamettu. Single-image vignetting correction using radial gradient symmetry. In *CVPR*, 2008.