

Single-Image Vignetting Correction Using Radial Gradient Symmetry

Yuanjie Zheng¹ Jingyi Yu¹ Sing Bing Kang² Stephen Lin³ Chandra Kambhamettu¹

¹University of Delaware, Newark, DE, USA {zheng, yu, chandra}@eecis.udel.edu

²Microsoft Research, Redmond, WA, USA SingBing.Kang@microsoft.com

³Microsoft Research Asia, Beijing, P.R. China stevelin@microsoft.com

Abstract

In this paper, we present a novel single-image vignetting method based on the symmetric distribution of the radial gradient (RG). The radial gradient is the image gradient along the radial direction with respect to the image center. We show that the RG distribution for natural images without vignetting is generally symmetric. However, this distribution is skewed by vignetting. We develop two variants of this technique, both of which remove vignetting by minimizing asymmetry of the RG distribution. Compared with prior approaches to single-image vignetting correction, our method does not require segmentation and the results are generally better. Experiments show our technique works for a wide range of images and it achieves a speed-up of 4-5 times compared with a state-of-the-art method.

1. Introduction

Vignetting refers to the intensity fall-off away from the image center, and is a prevalent artifact in photography. It is typically a result of the foreshortening of rays at oblique angles to the optical axis and obstruction of light by the stop or lens rim. This effect is sometimes deliberately added for artistic effects. Regardless, it is not desirable in computer vision applications that rely on reasonably precise intensity distributions for analysis. Such applications include shape from shading, image segmentation, and image mosaicing.

Various techniques have been proposed to determine the vignetting effect in an image. Some require specific scenes for calibration, which typically must be uniformly lit [2, 8, 20, 24]. Others use image sequences with overlapping views [5, 13] or image sequences captured with a projector at different exposures and different aperture settings [7]. A more flexible technique was proposed by Zheng *et al.* [25]; it requires only a single (almost arbitrary) image.

Single image-based vignetting correction is more convenient in practice, especially when we have access to only one image and the camera source is unknown (as is typically the case for images lifted from the web). The chal-

lenge is to differentiate the global intensity variation of vignetting from those caused by local texture and lighting. Zheng *et al.* [25] treat intensity variation caused by texture as “noise”; as such, they require some form of robust outlier rejection in fitting the vignetting function. They also require segmentation and must explicitly account for local shading. All these are susceptible to errors.

We are also interested in vignetting correction using a single image. Our proposed approach is fundamentally different from Zheng *et al.*'s—we rely on the property of symmetry of the radial gradient distribution. (By radial gradient, we mean the gradient along the radial direction with respect to the image center.) We show that the radial gradient distribution for a large range of vignetting-free images is symmetric, and that vignetting always increases its skewness. We describe two variants for estimating the vignetting function based on minimizing the skewness of this distribution. One variant estimates the amount of vignetting at discrete radii by casting the problem as a sequence of least-squares estimations. The other variant fits a vignetting model using nonlinear optimization.

We believe our new technique is a significant improvement over Zheng *et al.* [25]. First, our technique implicitly accounts for textures that have no bearing in vignetting. It obviates the need for segmentation and, for one variant, requires fewer parameters to estimate. In addition to the better performance, our technique runs faster, from 4-5 minutes [25] to less than 1 minute for a 450×600 image in a 2.39 GHz PC.

2. Natural Image Statistics

Our method assumes that the distributions of radial gradients in natural images are statistically symmetric. In this section, we first review the distribution properties of image gradients and confirm the validity of our assumption. We then show the effect of vignetting on the gradient distribution.

2.1. Symmetry of Image Gradients

Recent research in natural image statistics has shown that images of real-world scenes obey a heavy-tailed distribution in their gradients: it has most of its mass on small values but gives significantly more probability to large values than a Gaussian [4, 26, 11]. If we assume image noise to be negligible, a distribution of radial gradients $\psi(I)$ will have a similar shape, as exemplified in Fig. 2 (b). $\psi(I)$ is also highly symmetric around the distribution peak, especially among small gradient magnitudes. This characteristic arises from the relatively small and uniform gradients (e.g., textures) commonly present throughout natural images. $\psi(I)$ is generally less symmetric near the tails, which typically represent abrupt changes across shadow and occlusion boundaries and tend to be less statistically balanced. Furthermore, recent work [15] has shown that it is reasonable to assume image noise to be symmetric when the radiometric camera response is linear. This implies that including noise in our analysis will not affect the symmetry of the gradient distribution.

The symmetric, heavy-tailed shape characteristics of the gradient distributions have been exploited for image denoising, deblurring, and superresolution [18, 19, 12, 21, 10, 3, 1, 11, 22]. Fergus *et al.* [3] and Weiss *et al.* [23] used a zero-mean mixture-of-Gaussians to model the distributions of horizontal and vertical gradients for image deblurring. Huang *et al.* [6] use a generalized Laplacian function based on the absolute values of derivatives. Roth *et al.* [18] apply the Student's t-distribution to model this distribution for image denoising. Levin *et al.* [11] fit the distribution with an exponential function of the gradient magnitude. Zhu *et al.* [26] choose a Gibbs function in which the potential function is an algebraic expression of the gradient magnitude.

2.2. Radial Gradient

In this paper, we study the distribution of a special type of gradient, the radial gradient (RG). The radial gradient is the image gradient along the radial direction with respect to the image center, as shown in Fig. 1. With the optical center at (x_0, y_0) , the radial gradient at each pixel (x, y) can be computed as

$$\psi_r^I(x, y) = \begin{cases} |\nabla I(x, y) \cdot \vec{r}(x, y)| / |\vec{r}(x, y)| & |\vec{r}(x, y)| > 0 \\ 0 & |\vec{r}(x, y)| = 0 \end{cases} \quad (1)$$

where

$$\nabla I(x, y) = \left[\frac{\partial I}{\partial x}, \frac{\partial I}{\partial y} \right]^T, \quad \vec{r}(x, y) = [x - x_0, y - y_0]^T.$$

As with the horizontal and vertical gradients, the radial gradient distribution (which we call the RG distribution) in a vignetting-free image is also near-symmetric and heavy-




Figure 1. Demonstration for the definition of radial gradient.

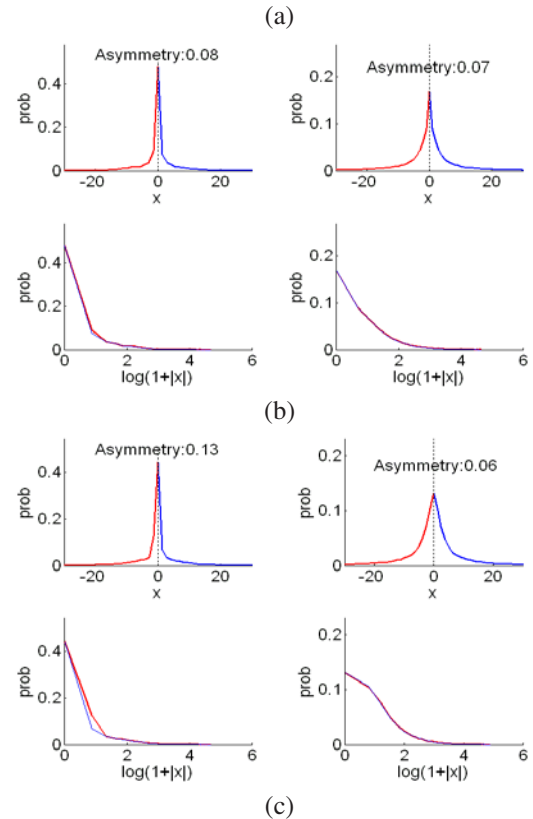


Figure 2. Gradient histograms for two natural images (a). In (b) and (c), top to bottom: regular histogram and corresponding $\log(1 + |x|)$ histogram. (b) are plots for horizontal gradients while (c) are for radial gradients.

tailed, as shown in Fig. 2. On the other hand, the RG distribution of an image with vignetting is asymmetric or skewed, as shown at the bottom left in Fig. 2(c).

We show both the regular and $\log(1 + |x|)$ histograms. In the regular histogram, x is the gradient value while “prob” denotes its density. The $\log(1 + |x|)$ histogram (e.g., in Fig. 2) is obtained by mapping x to $\log(1 + |x|)$. This mapping enhances any asymmetry that is present near the peak of the histogram. Note that the curve for negative x is

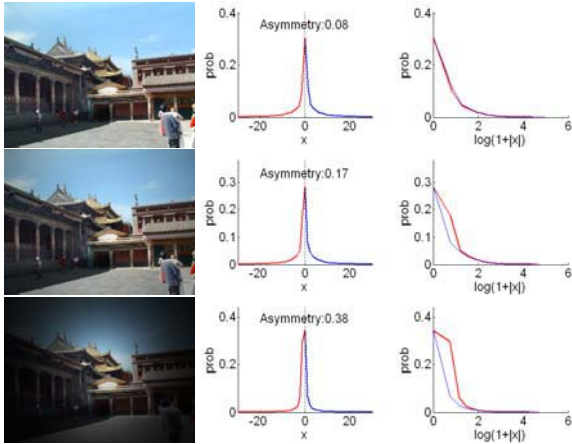


Figure 3. Comparison of skewness of RG distributions for varying degrees of vignetting. From left to right: image, histogram of radial gradients and skewness (asymmetry measure), and $\log(1 + |x|)$ histogram. From top to bottom: increasing degrees of vignetting.

folded over to the positive side (hence the two curves, with red representing negative x 's and blue representing positive x 's). Section 3.1 describes how we measure the skewness of the gradient histogram distribution.

Since vignetting is radial in nature, it is convenient to analyze in polar coordinates:

$$Z(r, \theta) = I(r, \theta)V(r), \quad (2)$$

where Z is the image with vignetting, I is the vignetting-free image, and V is the vignetting function. (The coordinate center corresponds to the image center.) Notice that V is a function of r only; this is because it can be assumed to be rotationally symmetric [2, 8, 20, 24, 25].

The radial gradient in polar coordinates is then computed as

$$\frac{dZ(r, \theta)}{dr} = \frac{dI(r, \theta)}{dr}V(r) + I(r, \theta)\frac{dV(r)}{dr}. \quad (3)$$

Let us now consider the RHS of equation (3). The first term simply scales the radial gradients by V . Since V is radially symmetric, the scaled distribution of the first term is expected to be mostly symmetric for natural images. The distribution of the second term, however, is not. This is because vignetting functions are radially monotonically decreasing, i.e., $\frac{dV(r)}{dr} \leq 0$. Since the scene radiance I is always positive, the second term is always negative. Therefore, the distribution of the second term is asymmetric. Furthermore, the more severe the vignetting, the more asymmetric the RG distribution of Z will be, as shown in Fig. 3. Moreover, with the same vignetting function, brighter scenes with larger I will exhibit greater asymmetry in the distribution of the second term. This is consistent with the common observation that vignetting is more obvious in a brighter scene, as shown in Fig. 4.

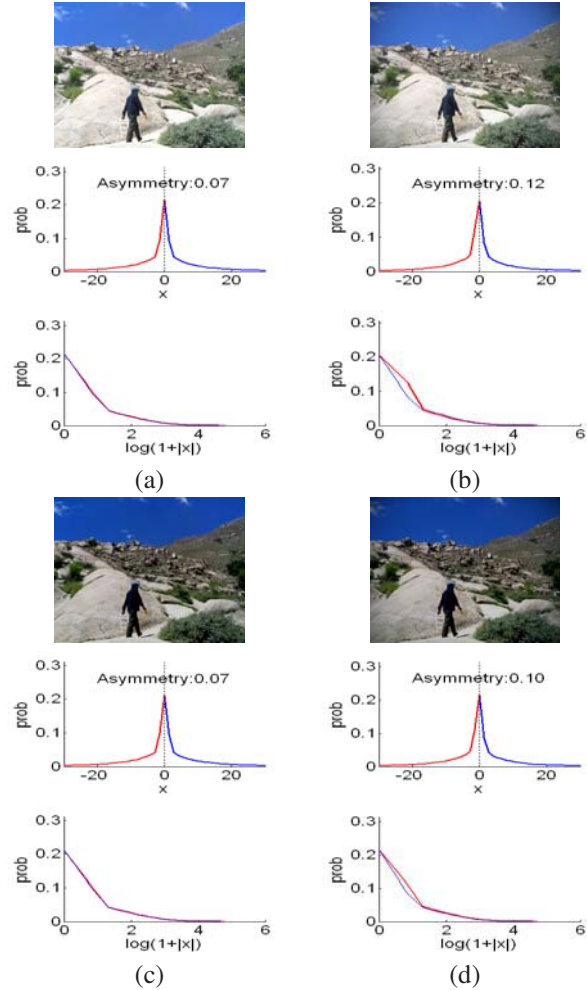


Figure 4. Effect of darker images on skewness. (a) Original image, (b) image with synthetic vignetting, (c) darkened version of (a), (d) same amount of synthetic vignetting applied to (c). For each of (a)-(d), from top to bottom: image, histogram, $\log(1 + |x|)$ histogram. Notice that brighter images with vignetting has a greater skewness.

In contrast to radial gradients, the symmetry of horizontal and vertical gradient distributions in an image is relatively unaffected by vignetting. Since vignetting is radially symmetric from the image center, it can be seen as increasing the magnitudes of horizontal or vertical gradients on one side of the image, while decreasing the gradient magnitudes on the other side. The vignetting-free gradient distributions of each side of the image can be assumed to be symmetric, and increasing or decreasing their magnitudes will in general leave the distributions symmetric. As a result, horizontal and vertical gradient distributions do not provide the vignetting information that is available from radial gradients.

3. Vignetting Estimation with Radial Gradients

In this section, we describe two variants of our single image vignetting correction technique based on minimizing the asymmetry of the RG distribution. One variant estimates the amount of vignetting at discrete radii by casting the problem as a sequence of least-squares optimizations. The other variant fits an empirical vignetting model by non-linear optimization.

3.1. Asymmetry Measure

We start by describing our quantitative measure of distribution function asymmetry. We use the Kullback-Leibler (K-L) divergence that describes the relative entropy between the two sides of a distribution. Let $H(\psi)$ be the histogram of gradient ψ centered at zero radial gradient. We compute the positive and negative sides of the RG distribution as

$$H_+(\psi) = \begin{cases} \frac{1}{A_1}H(\psi) & \psi \geq 0 \\ 0 & \psi < 0 \end{cases}, \quad (4)$$

$$H_-(\psi) = \begin{cases} \frac{1}{A_2}H(-\psi) & \psi \geq 0 \\ 0 & \psi < 0 \end{cases}, \quad (5)$$

where A_1 and A_2 are normalization factors that map the histograms to probability distribution functions. They are defined as

$$A_1 = \sum_{\psi \geq 0} H(\psi), \quad A_2 = \sum_{\psi \leq 0} H(\psi). \quad (6)$$

The K-L divergence measures the difference between probability distributions $H_+(\psi)$ and $H_-(\psi)$ as

$$\sum_{\psi} \left(H_+(\psi) \cdot \log \frac{H_+(\psi)}{H_-(\psi)} \right). \quad (7)$$

Note that two different histograms may still correspond to two similar probability distributions after normalization. We account for this difference by incorporating the normalization factors in our asymmetry measure Γ :

$$\Gamma(I) = \lambda_h \sum_{\psi} \left(H_+(\psi^I) \cdot \log \frac{H_+(\psi^I)}{H_-(\psi^I)} \right) + (1 - \lambda_h) |A_1 - A_2|^{\frac{1}{4}}. \quad (8)$$

This asymmetry measure is applied to both horizontal and radial gradient distributions. In this paper, we use $\Gamma_r(I)$ and $\Gamma_h(I)$ to represent the asymmetry measure of the RG distribution and horizontal gradient distribution of image I , respectively.

We have compared Γ_r with Γ_h on images in the Berkeley Segmentation Dataset [14] and found Γ_r to be considerably more sensitive to vignetting. For this dataset, Γ_r is significantly higher on average than Γ_h (0.12 vs. 0.08). In Fig. 5,



Figure 5. Images (from the Berkeley Segmentation Dataset) sorted by asymmetry. The top row images have the highest asymmetry while the bottom row images have the lowest.

we display in the top row the four images with the highest Γ_r . The bottom row shows the four images with the lowest Γ_r . Vignetting is clearly strong in the top four images, while the bottom four are practically vignetting-free.

We have also compared Γ_r and Γ_h before and after vignetting correction by the method in [25]. With vignetting correction, significant reductions in Γ_r were observed, from an average of 0.12 down to 0.072 over 40 images. In contrast, no obvious changes were observed for Γ_h (0.074 vs. 0.076). Note that vignetting correction brings Γ_r down to a level similar to that of Γ_h (0.072 vs. 0.076). We repeated these vignetting correction experiments on log intensity images and found that their RG and horizontal gradient distributions also follow these trends.

Based on this asymmetry measure, we propose two variants for minimizing skewness: (1) a least-squares solution with discrete radii, and (2) a nonlinear model-based solution.

3.2. Least-squares Solution with Discrete Radii

Our goal is to find the optimal vignetting function V that minimizes asymmetry of the RG distribution. By taking the log of equation (2), we get

$$\ln Z(r, \theta) = \ln I(r, \theta) + \ln V(r). \quad (9)$$

Let $\mathcal{Z} = \ln Z$, $\mathcal{I} = \ln I$, and $\mathcal{V} = \ln V$. We denote the radial gradients of \mathcal{Z} , \mathcal{I} , and \mathcal{V} for each pixel (r, θ) by $\psi_r^{\mathcal{Z}}(r, \theta)$, $\psi_r^{\mathcal{I}}(r, \theta)$, and $\psi_r^{\mathcal{V}}(r, \theta)$, respectively. Then,

$$\psi_r^{\mathcal{Z}}(r, \theta) = \psi_r^{\mathcal{I}}(r, \theta) + \psi_r^{\mathcal{V}}(r). \quad (10)$$

Given an image \mathcal{Z} with vignetting, we find a maximum a posteriori (MAP) solution to \mathcal{V} . Using Bayes' rule, this amounts to solving the optimization problem

$$\mathcal{V} = \arg \max_{\mathcal{V}} P(\mathcal{V}|\mathcal{Z}) \propto \arg \max_{\mathcal{V}} P(\mathcal{Z}|\mathcal{V})P(\mathcal{V}). \quad (11)$$

We consider the vignetting function at discrete, evenly-sampled radii: $(V(r_t), r_t \in S_r)$, where $S_r = \{r_0, r_1, \dots, r_{n-1}\}$. We also partition an image into sectors divided along these discrete radii, such that r_m is the inner

radius of sector m . Each pixel (r, θ) is associated with the sector in which it resides, and we denote sector width by δr .

The vignetting function is in general smooth; therefore, we impose a smoothness prior over \mathcal{V} :

$$P(\mathcal{V}) = e^{-\lambda_s \sum_{r_t \in S_r} \mathcal{V}''(r_t)^2}, \quad (12)$$

where λ_s is chosen to compensate for the noise level in the image, and $\mathcal{V}''(r_t)$ is approximated as

$$\mathcal{V}''(r_t) = \frac{\mathcal{V}(r_{t-1}) - 2\mathcal{V}(r_t) + \mathcal{V}(r_{t+1}))}{(\delta r)^2}.$$

To compute $P(\mathcal{Z}|\mathcal{V})$, from equation (10), we have

$$\psi_r^{\mathcal{I}}(r, \theta) = \psi_r^{\mathcal{Z}}(r, \theta) - \psi_r^{\mathcal{V}}(r). \quad (13)$$

We impose the sparsity prior [11, 9] on the vignetting-free image \mathcal{I} :

$$P(\mathcal{Z}|\mathcal{V}) = P(\psi_r^{\mathcal{I}}) = e^{-|\psi_r^{\mathcal{I}}|^\alpha}, \quad \alpha < 1. \quad (14)$$

$|\psi_r^{\mathcal{I}}|$ is used because of symmetry of the RG distribution for \mathcal{I} .

Substituting equation (13) in equation (14), we have

$$P(\mathcal{Z}|\mathcal{V}) = e^{-\sum_{(r,\theta)} |\psi_r^{\mathcal{Z}}(r,\theta) - \psi_r^{\mathcal{V}}(r)|^\alpha}, \quad (15)$$

where $\psi_r^{\mathcal{V}}(r) = (\mathcal{V}(r_m) - \mathcal{V}(r_{m-1})) / \delta r$, with m denoting the sector within which the pixel (r, θ) resides. The overall energy function $P(\mathcal{Z}|\mathcal{V})P(\mathcal{V})$ can then be written as

$$\mathcal{O} = \sum_{(r,\theta)} |\psi_r^{\mathcal{Z}}(r,\theta) - \psi_r^{\mathcal{V}}(r)|^\alpha + \lambda_s \sum_{r_t \in S_r} \mathcal{V}''(r_t)^2. \quad (16)$$

Our goal is to find the values of $V(r_t)$, $t = \{0, 1, \dots, n-1\}$ that minimize \mathcal{O} .

To effectively apply this energy function, a proper sparsity parameter α for the RG distribution of \mathcal{I} must be selected. As given in equation (14), α must be less than 1. However, very small values of α allow noise to more strongly bias the solution [26, 11]. We have empirically found that values of α between 0.3 and 0.9 yield robust estimates of the vignetting function for most images. For $0 < \alpha < 1$ though, equation (16) does not have a closed-form solution.

To optimize equation (16), we employ an iteratively reweighted least squares (IRLS) technique [9, 16]. IRLS poses the optimization as a sequence of standard least squares problems, each using a weight factor based on the solution of the previous iteration. Specifically, at the k th iteration, the energy function using the new weight can be written as

$$\mathcal{O}_k = \sum_{(r,\theta)} w_k(r, \theta) (\psi_r^{\mathcal{Z}}(r, \theta) - \psi_r^{\mathcal{V}_k}(r))^2 + \lambda_s \sum_{r_t \in S_r} \mathcal{V}_k''(r_t)^2. \quad (17)$$




Figure 6. Computed weights (equation (17)) in the least squares variant after the 3rd iteration of the IRLS algorithm.

The weight $w_k(r, \theta)$ is computed in terms of the optimal \mathcal{V}_{k-1} from the last iteration as

$$\begin{aligned} w_k(r, \theta) &= e^{-S_1} (1 - e^{-S_2}), \\ S_1 &= \left| \psi_r^{\mathcal{Z}}(r, \theta) - \psi_r^{\mathcal{V}_{k-1}}(r) \right|, \\ S_2 &= \alpha (S_1)^{\alpha-1}. \end{aligned} \quad (18)$$

The energy function then becomes a standard least-squares problem, which allows us to optimize \mathcal{V}_k using SVD.

In our experiments, we initialized $w_0(i, j) = 1$ for all pixels (i, j) , and found that it suffices to iterate 3 or 4 times to obtain satisfactory results. We also observed that the re-computed weights at each iteration k are higher at pixels whose radial gradients in \mathcal{Z} are more similar to the ones in the estimated \mathcal{V}_{k-1} . Thus, the solution is biased towards smoother regions whose radial gradients are relatively smaller. In addition, in a departure from [9], the recomputed weights in our problem are always within the range $[0, 1]$. Fig. 6 shows the weights recovered at the final iteration for an indoor image.

Our IRLS approach for estimating the vignetting function does not require any prior on the vignetting model. However, it requires choosing a proper coefficient λ_s to balance the smoothness prior on V and the radial gradient prior on \mathcal{I} . Since we choose a relatively small value of α , our vignetting estimation is biased more towards smooth regions than sharp edges. In essence, we emphasize the central symmetric part of the RG distribution rather than the less symmetric heavy tails.

The IRLS variant has the advantage of fast convergence and a linear solution. However, it requires estimating many parameters, each corresponding to a discrete radius value. We now describe the second variant, which is model-based and requires far fewer number of parameters to estimate.

3.3. Model-based Solution

Many vignetting models exist, including polynomial functions [2, 20], hyperbolic cosine functions [24], as well as physical models that account for the optical and geometrical causes of vignetting such as off-axis illumination and light path obstruction [2, 8]. In this paper, we use the extended Kang-Weiss model [25] in which brightness ratios are described in terms of an off-axis illumination factor A ,

a geometric factor G (represented by a polynomial), and a tilt factor. By neglecting the tilt factor, we have

$$V(r) = A(r)G(r), \quad (r) \in \Omega \quad (19)$$

$$A(r) = \frac{1}{(1 + (r/f)^2)^2},$$

$$G(r) = (1 - \alpha_1 r - \dots - \alpha_p r^p),$$

where f is the effective focal length of the camera and a_1, \dots, a_p are the coefficients of the p th order polynomial associated with G . In our experiments, $p = 5$.

We estimate the parameters in this vignetting model, i.e., f, a_1, \dots, a_p , by minimizing

$$\mathcal{O} = \lambda \Gamma_r \left(\frac{Z}{V} \right) + (1 - \lambda) \left(\frac{N_b}{N_\Omega} \right)^{1/4}, \quad (20)$$

where $\Gamma_r \left(\frac{Z}{V} \right)$ is the measure of asymmetry for image Z/V using equation (8), N_Ω is the total number of pixels in the image, and N_b is the number of pixels whose estimated vignetting values lie outside the valid range $[0, 1]$ or whose corrected intensities exist outside of $[0, 255]$. In essence, the second term in equation (20) penalizes outlier pixels.

To find the optimal vignetting model, we minimize the energy function in (20) using the Levenberg-Marquardt (L-M) algorithm [17]. We first solve for the focal length by fixing the geometric factor G to be 0. We then fix focal length and compute the optimal coefficients a_1, \dots, a_p of the geometric factor. Finally, we use the estimated focal length and geometric coefficients as an initial condition and re-optimize all parameters using the L-M method.

There are many advantages of using the vignetting model in equation (19). First, it effectively models the off-axis illumination effect $A(r)$ using a single parameter f . The off-axis illumination effect accounts for a prominent part of the vignetting for natural images. Second, as shown in Fig. 7, the profile of the energy function (20) with respect to focal length enables quick convergence by L-M optimization when estimating the focal length. Finally, the polynomial parameters in the extended Kang-Weiss model can effectively characterize the residual vignetting effect after removing the off-axis effect. In our experiments, by initializing these parameters simply to 0, the L-M method can quickly converge to satisfactory solutions.

4. Results

We applied our algorithms on images captured using a Canon G3, Canon EOS 20D, and Nikon E775, as well as on images from the Berkeley Segmentation Dataset [14]. The top row in Fig. 5 show four images from the Berkeley Database with the strongest degree of vignetting. We apply our least-squares and model-fitting methods to these images, and as seen in Fig. 8, the results are good.




Figure 7. Model-based vignetting correction. (a) Input image, (b) final corrected image, and (c) graph of objective function (20) vs. focal length. The images above the graph, from left to right, correspond to corrected versions using focal length values indicated by green squares on the curve. The focal length yielding the minimum value is the final solution.

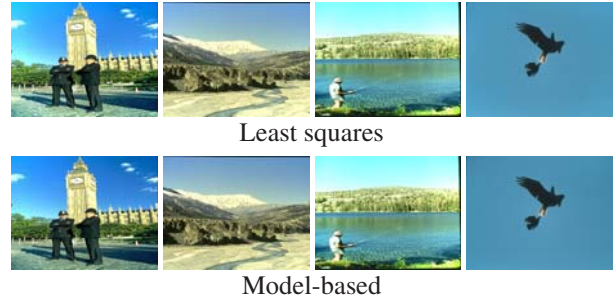


Figure 8. Vignetting correction results using our methods on the four most heavily vignettted images in the Berkeley Segmentation Dataset (Fig. 5).

We ran our algorithms on 20 indoor images. The vignetting artifacts in indoor images are generally difficult to correct due to greater illumination non-uniformity [25]. Since our methods are based on modeling the asymmetry of the gradient distributions instead of the intensity distributions, they are robust in vignetting estimation for indoor images. The results shown in the top rows of Fig. 9 demonstrate that our methods are able to effectively reduce vignetting despite highly non-uniform illumination.

We have also tested our methods on 15 highly textured images. While many previous approaches rely on robust segmentation of textured regions, our methods uniformly model the more slowly-varying vignetting and the high-frequency textures in terms of the radial gradient distributions: the textures correspond to the heavy tails of the dis-

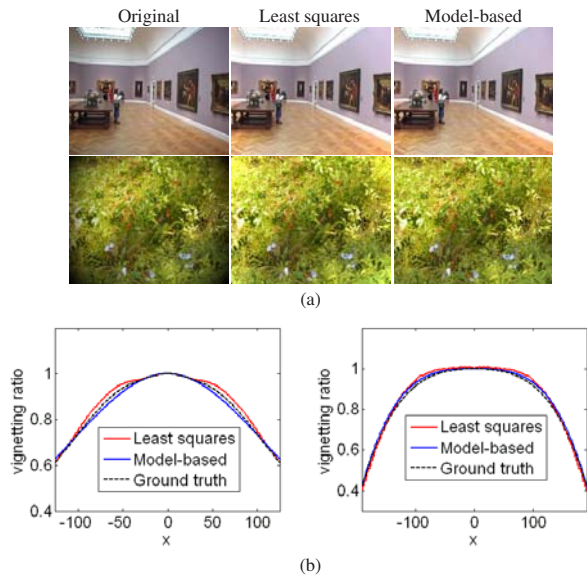


Figure 9. Results on indoor and textured images. (a) From left to right: input image, corrected image using least squares, corrected image using the model-based variant. (b) From left to right: estimated vignetting curves for images in (a). The red curves are obtained by least squares, the blue curves are obtained by the model-based method, and the black dotted curves are the ground truth.

	Zheng <i>et al.</i>	Least squares	Model-based
Time	285 sec	35 sec	51 sec

Table 1. Comparison of average execution time on 70 images.

tribution and vignetting is reflected in the asymmetry of the distribution. Therefore, without segmentation, our methods can still significantly reduce vignetting in the presence of strong textures, such as leaves on a tree, as shown in the bottom row of Fig. 9.

We have compared the speed between our methods and the previous single-image vignetting correction method [25] on a total of 70 outdoor, indoor, and textured images. All images have a resolution of 450×600 and all algorithms were implemented in Matlab (except for the segmentation component of [25] in C++) and run on a Dell PC with 2.39 GHz Intel Core 2 CPU. Our algorithms achieved on average a speed-up of 4-5 times compared with Zheng *et al.*'s algorithm (see Table 1). This is mainly because our methods do not require iterative segmentation and vignetting correction.

To evaluate accuracy, we obtained ground truth vignetting functions using an approach similar to that described in [25]: we captured multiple images of a distant white surface under approximately uniform illumination. Table 2 lists residual errors for our methods as well as Zheng *et al.*'s algorithm [25]. For outdoor scenes, our model-fitting variant performs the best while the method of Zheng *et al.* and our least-squares variant are comparable. For indoor and texture scenes, our two methods, in

	Zheng <i>et al.</i>	Least squares	Model-based
Outdoor	1.9/0.5	1.9/1.0	1.4/0.3
Indoor	2.9/1.8	2.4/1.3	2.5/1.2
Texture	5.7/2.1	5.3/2.4	4.0/1.9

Table 2. Comparison of mean/standard-deviation of the Mean Squared Errors ($\times 10^{-3}$) for 70 images.

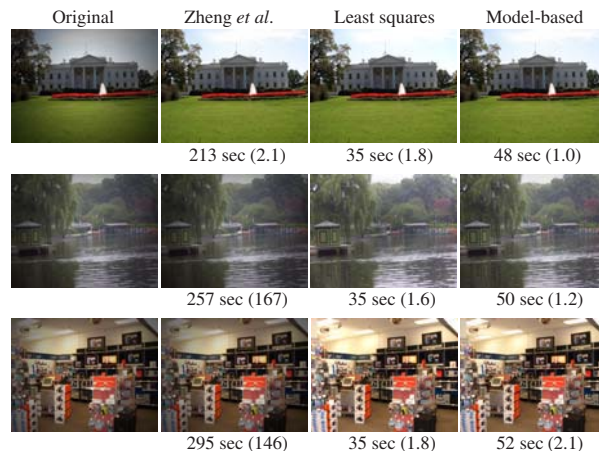


Figure 10. Comparisons of speed and accuracy. The numbers within parentheses are mean squared errors ($\times 10^{-3}$).



Figure 11. Final segmentations on the images in Fig. 10 by the vignetting correction method of Zheng *et al.*

particular the model-based method, estimate the vignetting functions more accurately. This is mainly because our technique is based on the symmetry of the RG distribution while the method by Zheng *et al.* [25] relies on the (less reliable) measurement of homogeneity in textures and colors. RG symmetry holds for a wide range of natural images even though they contain few homogeneous regions (e.g., highly textured images). It is thus not surprising that our methods are able to correct vignetting in images with highly complex textures or non-uniform illumination while the method of Zheng *et al.* is less able to, as shown in Fig. 10. Fig. 11 exemplifies the problem of using segmentation for vignetting removal. Notice that many of the segments in the second and third images cover regions that are either non-uniformly textured or are inhomogeneous, resulting in sub-optimal results.

5. Discussion

Our model-based vignetting variant uses a small number of parameters, and as such, has a better chance of converg-

ing to an optimal solution. However, since its optimization is nonlinear, convergence is slower than the least squares variant. Unfortunately, not all images with vignetting fit the Kang-Weiss vignetting model. Cameras with specially designed lenses, for example, may produce vignetting functions that deviate from this model. Here, the more flexible least squares variant would perform better.

A major limitation of our techniques is the assumption of the optical center being at the image center. Our techniques would not work for images cropped off-center. While it is possible to search for the optical center, issues of convergence would have to be dealt with effectively.

6. Conclusion

We have presented a novel single-image vignetting correction method based on the symmetric distribution of the radial gradient (RG). The radial gradient is the image gradient along the radial direction with respect to the image center. We have shown for natural images without vignetting that the RG distribution is generally symmetric, while it will be skewed if the image is corrupted by vignetting. To remove vignetting, we have developed two variants for correcting the asymmetry of the RG distribution. One variant estimates the amount of vignetting at discrete radii by casting the problem as a sequence of least-squares estimations. The other variant fits a vignetting model using nonlinear optimization.

Our techniques avoid the segmentation that is required by previous methods. Instead, we model the symmetry of the RG distribution over the entire image. Experiments on a wide range of natural images have shown that our techniques are overall more robust and accurate, particularly for images with textures and non-uniform illuminations. These images are difficult to handle effectively using segmentation-based approaches. Our methods are also faster than the segmentation-based approaches. Both methods achieve a speed-up of 4-5 times compared with a state-of-the-art method, and with comparable or better results.

References

- [1] N. Apostoloff and A. Fitzgibbon. Bayesian video matting using learnt image priors. In *CVPR*, 2004. 2
- [2] N. Asada, A. Amano, and M. Baba. Photometric calibration of zoom lens systems. In *Proc. Int. Conf. on Pattern Recognition*, pages 186–190, 1996. 1, 3, 5
- [3] R. Fergus, B. Singh, A. Hertzmann, S. T. Roweis, and W. T. Freeman. Removing camera shake from a single photograph. *ACM Transactions on Graphics*, 25(3):787–794, 2006. 2
- [4] D. J. Field. What is the goal of sensory coding? *Neural Computation*, 6(4):559–601, 1994. 2
- [5] D. Goldman and J. Chen. Vignette and exposure calibration and compensation. In *ICCV*, pages 899–906, 2005. 1
- [6] J. Huang and D. Mumford. Statistics of natural images and models. In *ICCV*, 1999. 2
- [7] R. Juang and A. Majumder. Photometric self-calibration of a projector-camera system. In *CVPR*, 2007. 1
- [8] S. Kang and R. Weiss. Can we calibrate a camera using an image of a flat textureless lambertian surface? In *European Conf. on Computer Vision*, volume II, pages 640–653, 2000. 1, 3, 5
- [9] A. Levin, R. Fergus, F. Durand, and W. T. Freeman. Image and depth from a conventional camera with a coded aperture. *ACM Trans. on Graphics*, 26(3), 2007. 5
- [10] A. Levin and Y. Weiss. User assisted separation of reflections from a single image using a sparsity prior. In *ECCV*, 2004. 2
- [11] A. Levin, A. Zomet, and Y. Weiss. Learning to perceive transparency from the statistics of natural scenes. In *NIPS*, 2002. 2, 5
- [12] A. Levin, A. Zomet, and Y. Weiss. Learning how to inpaint from global image statistics. In *ICCV*, volume 1, pages 305–312, 2003. 2
- [13] A. Litvinov and Y. Schechner. Addressing radiometric non-idealities: A unified framework. In *CVPR*, pages 52–59, 2005. 1
- [14] 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*, volume 2, pages 416–423, July 2001. 4, 6
- [15] Y. Matsushita and S. Lin. Radiometric calibration from noise distributions. In *CVPR*, 2007. 2
- [16] P. Meer. *Robust techniques for computer vision*, pages 107–190. Prentice-Hall, 2005. 5
- [17] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, New York, NY, USA, 1992. 6
- [18] S. Roth and M. Black. Fields of experts: A framework for learning image priors. In *CVPR*, pages 860–867, 2005. 2
- [19] S. Roth and M. J. Black. Steerable random fields. In *ICCV*, 2007. 2
- [20] A. A. Sawchuk. Real-time correction of intensity nonlinearities in imaging systems. *IEEE Trans. on Computers*, 26(1):34–39, 1977. 1, 3, 5
- [21] M. Tappen, B. Russell, and W. Freeman. Exploiting the sparse derivative prior for super-resolution and image demosaicing. In *IEEE Workshop on Statistical and Computational Theories of Vision*, 2003. 2
- [22] Y. Weiss. Deriving intrinsic images from image sequences. In *ICCV*, 2001. 2
- [23] Y. Weiss and W. T. Freeman. What makes a good model of natural images? In *CVPR*, 2007. 2
- [24] W. Yu. Practical anti-vignetting methods for digital cameras. *IEEE Trans. on Cons. Elect.*, 50:975–983, 2004. 1, 3, 5
- [25] Y. Zheng, S. Lin, and S. B. Kang. Single-image vignetting correction. In *CVPR*, 2006. 1, 3, 4, 5, 6, 7
- [26] S. C. Zhu and D. Mumford. Prior learning and gibbs reaction-diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(11):1236–1250, 1997. 2, 5