



Published in final edited form as:

IEEE Trans Pattern Anal Mach Intell. 2013 June ; 35(6): 1480–1494. doi:10.1109/TPAMI.2012.210.

Single-Image Vignetting Correction from Gradient Distribution Symmetries

Yuanjie Zheng [Member, IEEE],

Department of Radiology, University of Pennsylvania, Philadelphia, PA 19104.

Yuanjie.Zheng@uphs.upenn.edu

Stephen Lin [Member, IEEE],

Microsoft Research Asia, No. 5 Danling Street, Haidian District, Beijing, China 100080.

stevlin@microsoft.com.

Sing Bing Kang [Fellow, IEEE],

Microsoft Corporation, One Microsoft Way, Redmond, WA 98052. sbkang@microsoft.com.

Rui Xiao,

Department of Biostatistics and Epidemiology, University of Pennsylvania, Philadelphia, PA

19104. rxiao@mail.med.upenn.edu.

James C. Gee [Member, IEEE], and

Department of Radiology, University of Pennsylvania, Philadelphia, PA 19104.

James.Gee@uphs.upenn.edu.

Chandra Kambhamettu [Member, IEEE]

Department of Computer and Information Sciences, University of Delaware, Newark, DE

19716-2712. chandra@cis.udel.edu.

Abstract

We present novel techniques for single-image vignetting correction based on symmetries of two forms of image gradients: semicircular tangential gradients (SCTG) and radial gradients (RG). For a given image pixel, an SCTG is an image gradient along the tangential direction of a circle centered at the presumed optical center and passing through the pixel. An RG is an image gradient along the radial direction with respect to the optical center. We observe that the symmetry properties of SCTG and RG distributions are closely related to the vignetting in the image. Based on these symmetry properties we develop an automatic optical center estimation algorithm by minimizing the asymmetry of SCTG distributions, and also present two methods for vignetting estimation based on minimizing the asymmetry of RG distributions. In comparison to prior approaches to single-image vignetting correction, our methods do not rely on image segmentation and they produce more accurate results. Experiments show our techniques to work well for a wide range of images while achieving a speed-up of 3-5 times compared to a state-of-the-art method.

Keywords

Vignetting correction; camera calibration; low-level vision; bias correction; nonuniformity correction

1 Introduction

Vignetting refers to the radial fall-off of brightness away from the center of an image. It is mainly caused by foreshortening of light rays at oblique angles to the optical axis and by obstruction of light by the stop or lens rim. This effect is undesirable in computer vision applications that rely on accurate or consistent intensity measurements of the scene, such as shape from shading, image mosaicing, image-based rendering, stereo matching, optical flow estimation, and image segmentation. For these applications, correcting for the brightness distortions caused by vignetting is an important step towards obtaining high-quality results.

To correct for vignetting, the optical center generally needs to be known, since the intensity fall-off is commonly modeled as radially symmetric about that point. Most previous techniques on vignetting correction simply assume the optical center to be at the center of the image [11], [7], [33], [10]. However, in practice it may lie at a considerable distance from the image center [13], or the input image may have been cropped (with the original image not available). In such cases, the actual image center needs to be determined for accurate removal of vignetting effects.

In this paper, we address the problem of vignetting correction based on a single input image from an unknown camera. This problem setting makes our work more useful in practice and widely applicable than most previous techniques (which typically require an explicit calibration step or an image sequence as input). However, single-image vignetting correction is a difficult problem to solve mainly because of two challenges: differentiating global brightness variations caused by vignetting from those caused by textures or lighting, and accurately locating the optical center given just a single input image.

To address these issues, we introduce an approach that takes advantage of regularities in the statistical distributions of two types of image gradients, namely semicircular tangential gradients (SCTG) and radial gradients (RG). An SCTG is a gradient along the tangential direction of a circle centered at the optical center and passing through the pixel. We observe that the SCTG distribution for a broad range of vignetted images exhibits a more symmetric structure for more accurate estimates of the optical center. We find a similar property for the RG, which is a gradient along the radial direction with respect to the optical center. Specifically, the RG distribution of an image becomes more symmetric if it contains less vignetting.

From these gradient distribution properties, we develop a technique that identifies the optical center as the point which minimizes asymmetry in the SCTG distribution. Once the optical center is found, we recover a vignetting-free image by minimizing asymmetry in the RG distribution. Two variants are presented for determining the vignetting-free image. 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 parametric vignetting model using a nonlinear optimization process.

We believe that these techniques for single-image based vignetting correction bring significant improvements over the previous work on the same problem [43], [42], which are based on image segmentation. First, our presented methods are not sensitive to segmentation

quality and are less susceptible to brightness variations within regions caused by texture and lighting. Second, for parametric vignetting models, this approach requires much fewer parameters to estimate. Third, its accurate determination of the optical center leads to improved vignetting estimation. Finally, in addition to better performance, our techniques operate faster, from 5 minutes in [43] to less than 2 minutes for our methods on a 450×600 image in a 2.39 GHz PC.

2 Related Work

In this section, we briefly review work related to our approach: optical center estimation, vignetting estimation, and natural image statistics.

2.1 Optical Center Estimation

Various techniques have been proposed for optical center estimation. Some estimate the optical center by locating the center of an optical effect such as vignetting [18], radial lens distortion [36], vanishing point [34], or focus/defocus [37]. These techniques generally require specific calibration scenes or instruments, such as a uniform scene [18], a special calibration target [36], a cube [34], a high-frequency textured pattern [37], or a laser emitter [13]. Another approach to estimate the optical center is to use a general camera calibration procedure that also estimates other intrinsic or extrinsic camera parameters. Some of these methods require a particular calibration pattern or scene [32], [10], while others perform self-calibration using image sequences captured with fixed camera settings [4], [8]. In contrast, we estimate the optical center from a single (vignetted) image of an *arbitrary natural scene* without knowledge of the camera. This makes our technique highly practical.

2.2 Vignetting Estimation

Different techniques have been proposed to estimate the amount of vignetting in images. Some require specific scenes for calibration, which typically must be uniformly lit [2], [10], [28], [38], [24]. Others use image sequences with overlapping views [7], [11], [19], image sequences captured with a projector at different exposures and different aperture settings [9], or image sequences from a lens setting [12]. A more flexible technique was proposed by Zheng *et al.* [43], [42], which requires only a single (almost arbitrary) image.

The method of [43], [42] decomposes an image into regions that are homogeneous with respect to color or texture, and then estimates vignetting based on low-frequency brightness variations within each region. Outlier rejection is used to discount intra-region variations inconsistent with other regions or with a parametric vignetting model. It relies on accurate image segmentation and the optical center being at the image center.

The single-image vignetting correction techniques proposed in this paper are instead based on statistical distributions of the SCTG and RG over the image. They do not rely on image segmentation quality but rather they directly estimate the vignetting from global distribution data, which we will show to bring improvements in both performance and running time. This paper is an extension of the methods in [44], [41] in which we further characterize the theoretical space of solutions to the problem and include additional validation on the benefits of optical center estimation in vignetting correction. A recently published work [20]

inspired by the methods in [44], [41] accounts for statistical regularities in a series of general image derivatives. However, unlike with the SCTG and RG in this paper, it is unclear how vignetting affects the statistical distributions of general image derivatives.

2.3 Natural Image Statistics

Recent research on natural image statistics has shown that images of real-world scenes have gradient distributions that exhibit a certain shape: sharply-peaked at zero and heavy-tailed at large values. Such a distribution assigns higher probability to small gradient magnitudes [6], [45], [16]. Its shape also exhibits symmetry around its peak [25], especially among small gradient magnitudes because of the numerous small and offsetting gradients (e.g., textures) commonly present throughout natural images. The gradient distribution is generally less symmetric near the tails, which often represent abrupt changes across shadow and occlusion boundaries and tend to be less statistically balanced. This special shape of image gradient distribution is generally invariant to certain image processing operations such as image resizing, cropping, and contrast adjustment. It has been used in various applications such as image denoising, deblurring, superresolution and object recognition (e.g., [25], [27], [30], [5], [1], [16], [35], [31]).

Previous work [22] has shown that it is reasonable to assume that image noise is symmetric when the radiometric camera response is linear. This implies that including noise in our analysis is not likely to affect the symmetry of the gradient distribution.

3 Definitions

In this section, we provide formal definitions of the radial gradient (RG) and the semicircular tangential gradient (SCTG), upon which our estimation methods are based.

Let Z denote a given image with vignetting. To index a pixel in Z , we will interchangeably use its Euclidean coordinates (x, y) and polar coordinates (r, θ) for simplicity. These two types of coordinates are meaningful only when the origin (x_0, y_0) or (r_0, θ_0) of the coordinate system is specified. To facilitate vignetting analysis, we wish to place the origin at the optical center, since vignetting is typically radially symmetric about the optical center.

For a pixel (x, y) , the conventional image gradient $\nabla Z(x, y)$ is defined in the Euclidean coordinate system as

$$\nabla Z(x, y) = \left[\frac{\partial Z(x, y)}{\partial x}, \frac{\partial Z(x, y)}{\partial y} \right]^T \quad (1)$$

which is composed of the horizontal gradient (HG) and vertical gradient (VG), respectively.

An image gradient may instead be defined along any direction in the image, as the inner product of the conventional gradient vector in Eq. (1) and the unit vector in the given direction. Geometrically, this represents the projection of the conventional gradient onto the specified direction. In the following, different forms of gradients will be defined in this way.

3.1 Radial Gradients

The RG is defined as the image gradient along the radial direction with respect to the origin of the coordinate system, as shown in Fig. 1. It measures the projection of the conventional image gradient on the radial direction.

The radial direction at a pixel (x, y) is expressed by the following unit vector:

$$\bar{r}(x, y) = \begin{cases} \frac{[x-x_0, y-y_0]^T}{|[x-x_0, y-y_0]^T|} & |[x-x_0, y-y_0]^T| > 0 \\ [0, 0]^T & |[x-x_0, y-y_0]^T| = 0 \end{cases} \quad (2)$$

where $|\cdot|$ denotes the L_2 norm. In practice, Eq. (2) can be approximated by

$$\bar{r}(x, y) = \frac{[x-x_0, y-y_0]^T}{|[x-x_0, y-y_0]^T| + \epsilon}, \quad (3)$$

where ϵ is a very small value (e.g., 1×10^{-6}) to avoid division by zero. The RG magnitude is then expressed as

$$\psi_r^Z(x, y) = \nabla Z(x, y) \cdot \bar{r}(x, y), \quad (4)$$

where “ \cdot ” denotes the inner product operator.

3.2 Tangential Gradients

The conventional tangential gradient (TG) of a pixel is the image gradient along the tangential direction of the circle that is centered at the origin of the coordinate system and passes through this pixel, as illustrated in Fig. 2. It measures the projection of the conventional image gradient on the tangential direction.

The tangential direction of a pixel (x, y) can be obtained by rotating the radial direction by 90° counterclockwise:

$$\bar{t}(x, y) = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \bar{r}(x, y). \quad (5)$$

The TG magnitude is computed as

$$\psi_t^Z(x, y) = \nabla Z(x, y) \cdot \bar{t}(x, y). \quad (6)$$

In the polar coordinate system, it is expressed as

$$\psi_t^Z(r, \theta) = \frac{\partial Z(r, \theta)}{\partial \theta}. \quad (7)$$

3.3 Semicircular Tangential Gradients

For the purpose of optical center estimation, we introduce SCTG which is defined as the image gradient along either the tangential direction (rotating the radial direction by 90° counterclockwise) or its opposite direction (rotating the radial direction by 90° clockwise). We refer to these two cases as counterclockwise SCTG and clockwise SCTG, respectively. Both forms of SCTG can be written jointly as

$$\bar{s} t(x, y) = \text{sgn}(x, y) \bar{t}(x, y) \quad (8)$$

for which we set $\text{sgn}(x, y) = 1$ for counterclockwise SCTG and $\text{sgn}(x, y) = -1$ for clockwise SCTG. The SCTG magnitude is computed as

$$\psi_{st}^Z(x, y) = \nabla Z(x, y) \cdot \bar{s} t(x, y), \quad (9)$$

as shown in Fig. 2(a) for the counterclockwise SCTG.

When estimating the optical center, its true position (x_0, y_0) is unknown and therefore its estimate (x'_0, y'_0) is instead used as the origin of the coordinate systems in the above definitions. We denote a pixel in such coordinate systems by (x', y') or (r', θ') .

We assign to each pixel either the clockwise SCTG or the counterclockwise SCTG based on its position with respect to the line defined by (x'_0, y'_0) and (x_0, y_0) (shown in red in Fig. 2(b)). Here, let us consider (x_0, y_0) to be a hypothetical true optical center to which is we wish to update (x'_0, y'_0) . This line divides any circle centered at (x'_0, y'_0) into two parts (as shown in green and blue in Fig. 2(c)). Those pixels that lie on the green semicircle take the value of the clockwise SCTG, while those in blue are assigned the counterclockwise SCTG. Equivalently, the image is divided into two parts by this line, with the clockwise SCTG assigned to pixels on one side and the counterclockwise SCTG to the other side.

This strategy for assigning a clockwise or counterclockwise SCTG to a pixel (x', y') can be expressed analytically by setting $\text{sgn}(x', y')$ in Eq. (8) according to

$$\text{sgn}(x', y') = \begin{cases} 1, & \text{if } (0 \leq \theta_l < \pi) \text{ and } \left((0 \leq \theta' \leq \theta_l) \text{ or } (\theta_l + \pi \leq \theta' \leq 2\pi) \right) \\ 1, & \text{if } (\pi \leq \theta_l < 2\pi) \text{ and } (\theta_l - \pi \leq \theta' \leq \theta_l) \\ -1, & \text{otherwise,} \end{cases} \quad (10)$$

where θ_l is the rotation angle from the x direction to the ray from (x'_0, y'_0) to (x_0, y_0) (set to zero when $(x'_0, y'_0) = (x_0, y_0)$) and θ' is the polar angle of (x', y') with respect to (x'_0, y'_0) . With the SCTG value of each pixel determined according to Eqs. (5)-(10), different pixels may be assigned in different ways depending on the estimate of the optical center, as shown for different estimates in Fig. 2(c).

Note that this assignment strategy does not require the actual optical center (x_0, y_0) , but instead needs only the angular direction (θ_l in Eq. (10)) of (x_0, y_0) from (x'_0, y'_0) to define the line between them. The SCTG distribution over the image is computed using different

sampled values of these angles $\theta_l = \{0, \frac{\pi}{10}, \frac{2\pi}{10}, \dots, \pi\}$ and taking one that yields greatest asymmetry, asymmetry should generally be maximum for the correct value of θ_l as we will later describe in Sec. 4.1.

4 Gradient Distribution Properties

In this section, we introduce a quantitative measure on the asymmetry of a gradient histogram distribution, briefly describe certain properties of conventional gradient distributions, and explore a series of asymmetry properties on RG and SCTG distributions which are caused by vignetting. These asymmetry properties will later be used to estimate the vignetting in an image. Since the logarithm of image intensities has a gradient distribution similar to that of the original image [45], these gradient distribution properties generally hold for both the original image and its logarithm image.

To better visualize asymmetry properties, we plot the gradient distributions in various forms: as a regular histogram, a log-histogram, and a $\log(1 + |x|)$ histogram (e.g., Fig. 3). A regular histogram shows the probability density (denoted by “prob”) of the gradient (denoted by x). The log-histogram and $\log(1 + |x|)$ -histogram are obtained by mapping “prob” to its logarithmic value and by mapping x to $\log(1 + |x|)$, respectively. The log-histogram enhances asymmetry that is present at the two tails of the histogram, while the $\log(1 + |x|)$ -histogram emphasizes near any skewness the peak. Note that for plots of the $\log(1 + |x|)$ -histogram we fold the curve for negative x values over to the positive side (hence two curves, with red representing negative x and blue representing positive x).

4.1 Asymmetry Measure

We introduce a quantitative measure for the asymmetry of a gradient histogram distribution. In this measure, we use the Kullback-Leibler (K-L) divergence to describe the shape difference between the two sides of the distribution. Let $H(\psi)$ be the histogram of gradients ψ centered at zero. We compute the positive and negative sides of the distribution as

$$H_+(\psi) = \begin{cases} \frac{1}{A_1} H(\psi) & \psi \geq 0 \\ 0 & \psi < 0 \end{cases}, \quad (11)$$

$$H_-(\psi) = \begin{cases} \frac{1}{A_2} H(-\psi) & \psi \geq 0 \\ 0 & \psi < 0 \end{cases}, \quad (12)$$

where A_1 and A_2 are areas under the curve for the two sides of the distribution, which are respectively defined as

$$A_1 = \sum_{\psi \geq 0} H(\psi), \quad A_2 = \sum_{\psi \leq 0} H(\psi). \quad (13)$$

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

$$\sum_{\psi} \left(H_+(\psi) \cdot \log \frac{H_+(\psi)}{H_-(\psi)} \right). \quad (14)$$

Since two different sides of a distribution may become similar after normalization, we account for their difference by incorporating A_1 and A_2 in our asymmetry measure Γ :

$$\Gamma(I) = \lambda_h \sum_{\psi} \left(H_+(\psi^I) \cdot \log \frac{H_+(\psi^I)}{H_-(\psi^I)} \right) + (1 - \lambda_h) |A_1 - A_2|^q, \quad (15)$$

where the weight λ_h (balancing the two terms) and exponent q (adjusting the emphasis on a small or large area-under-curve deviation value) were set empirically to 0.7 and 0.25, respectively, in our experiments. We use $\Gamma_r(I)$, $\Gamma_{st}(I)$, $\Gamma_t(I)$, $\Gamma_h(I)$, and $\Gamma_v(I)$ to represent the asymmetry measure of the RG, SCTG, TG, HG, and VG, respectively.

4.2 Conventional Gradients

The histogram distributions of the conventional gradients, HG and VG, are generally smooth and exhibit a special shape. As described in Sec. 2.3 and shown in Fig. 3, this distribution shape is characterized by a sharp peak, heavy tails and symmetry. We observe that deviation from this structure occurs if a global brightness variation is added to the image. For example, increasing the brightness from left to right in an image will cause corresponding increases in HG values and lead to asymmetry in the HG histogram distribution. However, not all types of global brightness variation will alter this special histogram structure. For example, vignetting generally will not change this histogram structure because of the radial symmetry of vignetting.

We further observe that, for images without vignetting or other directional brightness variation, the image gradient distributions defined along directions other than the vertical or horizontal direction share the special structure shown by the conventional gradients. This can be seen from the distribution comparisons of HG, TG, and RG in Fig. 3. Similar findings are also presented in [31], [25]. As discussed in [25], the special shape holds for image gradients in any direction because images consist of mostly smooth areas and the structures in images tend to have starting and ending boundaries of similar contrast.

4.3 Radial Gradients

We observe several interesting properties of RG distributions. First, the RG distribution of an image with vignetting is asymmetric or skewed, as shown at the bottom left of Fig. 4(c). Second, the stronger the vignetting, the more asymmetric the RG distribution will be, as shown in Fig. 5. Finally, with the same vignetting function, brighter scenes will exhibit greater asymmetry, as shown in Fig. 6. These are validated through the following geometric analysis.

4.3.1 Geometric Analysis—In polar coordinates, a given image Z with vignetting is represented by

$$Z(r, \theta) = I(r, \theta) V(r), \quad (16)$$

where I is the vignetting-free image and V is the vignetting function. We assume V to be rotationally symmetric (a function of r only), as done in most previous work [2], [10], [28], [38], [43]. The RG is then expressed as

$$\frac{dZ(r, \theta)}{dr} = \frac{dI(r, \theta)}{dr} V(r) + I(r, \theta) \frac{dV(r)}{dr}. \quad (17)$$

Let us now consider the right-hand side of Eq. (17). 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 non-positive. The distribution of the second term has most of its mass on negative values and therefore is asymmetric. Furthermore, the second term becomes more negative for brighter scenes and stronger vignetting. Stronger vignetting produces a larger brightness attenuation and more negative values of $\frac{dV(r)}{dr}$, while brighter scenes have larger values of I in the second term. This results in greater asymmetry in the RG distribution as shown in Fig. 5 and Fig. 6. These findings are consistent with the common observations that vignetting is more visually prominent when it is stronger or the scene is brighter.

In contrast to the RG, the symmetry of an HG or VG distribution is relatively unaffected by vignetting. Since vignetting is radially symmetric from the image center, it can be seen as increasing the magnitudes of HG or VG 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, HG and VG distributions provide little information about vignetting in an image.

4.3.2 Experimental Validation—We have compared the asymmetry measure Γ_r for the RG with Γ_h for the HG on images in the Berkeley Segmentation Dataset [21] 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. 15, we display in the top row the four images with the highest Γ_r .

We have also compared Γ_r and Γ_h before and after vignetting correction by the method in [43]. With vignetting correction, significant reductions in Γ_r were observed, from an average of 0.12 down to 0.072 over 40 images. By 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 HG distributions also follow these trends.

4.4 Semicircular Tangential Gradients

The SCTG distribution has several interesting properties. First, the SCTG distribution exhibits the same special shape as conventional gradients when images are free of any global brightness variation. Second, this special distribution shape still exists for the SCTG when images contain vignetting and the origin of the coordinate system is located at the true optical center. Third, the SCTG distribution of vignetted images will be skewed if the origin of the coordinate system is displaced from the true optical center, as demonstrated in Fig. 7 and Fig. 8. The causes for this asymmetry in the SCTG distribution will be explained in the following geometric analysis.

4.4.1 Geometric Analysis—To geometrically analyze SCTG, we rewrite Eq. (16) as

$$Z(r', \theta') = I(r', \theta') V(r', \theta'), \quad (18)$$

where (r', θ') denotes polar coordinates with origin of the coordinate system at an estimated optical center (x'_0, y'_0) .

We perform the geometric analysis in the logarithm domain of the image. We denote $\ln Z$, $\ln I$, and $\ln V$ by \mathcal{Z} , \mathcal{I} , and \mathcal{V} , respectively, and represent the TGs of \mathcal{Z} , \mathcal{I} , and \mathcal{V} for each pixel (r', θ') by $\psi_t^{\mathcal{Z}}(r', \theta')$, $\psi_t^{\mathcal{I}}(r', \theta')$, and $\psi_t^{\mathcal{V}}(r', \theta')$ respectively. From Eq. (18), we have

$$\psi_t^{\mathcal{Z}}(r', \theta') = \psi_t^{\mathcal{I}}(r', \theta') + \psi_t^{\mathcal{V}}(r', \theta'). \quad (19)$$

By analyzing the right-hand side of Eq. (19), we obtain four special properties of the SCTG distribution, based on which we present an optical center estimation algorithm in Sec. 5.1.

Property 1: The SCTG distribution is symmetric when the estimated optical center is at the true optical center, i.e., $(x'_0, y'_0) = (x_0, y_0)$.

On the right-hand side of Eq. (19), the first term is the TG of the vignetting-free image \mathcal{I} , which is taken to have a symmetric distribution for natural images as examined in Sec. 4.2. The second term represents the TG for the vignetting component of the image. Vignetting is radially symmetric about the estimated optical center if it is co-located with the true optical center. Therefore, the tangential gradient $\psi_t^{\mathcal{V}}(r', \theta')$ would be zero, and the SCTG is also zero regardless of whether the SCTG is computed in the clockwise or counterclockwise direction. In short, the SCTG distribution is symmetric when the optical center is correctly estimated.

Property 2: The SCTG distribution is asymmetric when the estimated optical center is incorrect, i.e., $(x'_0, y'_0) \neq (x_0, y_0)$.

When the estimated optical center is displaced from the true optical center, as shown in Fig. 9(a), the tangential gradient $\psi_t^{\mathcal{Y}}(r', \theta')$ relative to the estimated optical center (x'_0, y'_0) is equal to the projection of the true radial gradient $\psi_r^{\mathcal{Y}}(r, \theta)$ at that point onto the semicircular tangent direction, expressed as

$$\psi_t^{\mathcal{Y}}(r', \theta') = \psi_r^{\mathcal{Y}}(r, \theta) \cos(\beta), \quad (20)$$

where β is the rotation angle from $\psi_t^{\mathcal{Y}}(r', \theta')$ to $\psi_r^{\mathcal{Y}}(r, \theta)$.

From the property that vignetting attenuation is radially non-decreasing away from the true optical center, we show that $\psi_t^{\mathcal{Y}}(r', \theta')$ is always non-negative. Let γ denote the rotation angle from $\psi_r^{\mathcal{Y}}(r, \theta)$ to the line from (r', θ') to (x'_0, y'_0) , as shown in Fig. 9. Since $\beta + \gamma = 90^\circ$, it can be seen by simple geometric reasoning on the triangle defined by (x_0, y_0) , (x'_0, y'_0) and (r', θ') that $0^\circ < \gamma < 180^\circ$ and $-90^\circ < \beta < 90^\circ$. This holds true both for pixels assigned clockwise or counterclockwise SCTG values. From Eq. (20) and the radially non-decreasing vignetting ($\psi_r^{\mathcal{Y}}(r, \theta) \geq 0$) we conclude that $\psi_t^{\mathcal{Y}}(r', \theta') \geq 0$. Since the first term on the right-hand side of Eq. (19) has a symmetric distribution about zero and the second term skews the distribution towards positive values, the SCTG distribution is asymmetric about zero for vignettted images in which the estimated optical center is incorrect.

Property 3: Asymmetry values of SCTG distributions generally increase with greater distances between the estimated and true optical centers.

As seen in Fig. 9, for each pixel that does not lie on the line through (x'_0, y'_0) and (x_0, y_0) , β decreases for increasing distances between (x'_0, y'_0) and (x_0, y_0) . For pixels that lie within the circle defined by the diameter between (x'_0, y'_0) and (x_0, y_0) , we have $90^\circ < \gamma < 180^\circ$ and a negative β . With negative values of β , the SCTG values for pixels within the circle become smaller (less positive) for increasing error in the estimated optical center. On the other hand, pixels outside the circle have $0^\circ < \gamma < 90^\circ$ and a positive β , which leads to larger (more positive) SCTG values. Consequently from Eq. (20), $\psi_t^{\mathcal{Y}}(r', \theta')$ becomes smaller for pixels inside the circle but larger for pixels outside the circle. Since there is typically a much larger number of pixels outside the circle than inside, the SCTG distribution will become more skewed in the positive direction when the distance between (x'_0, y'_0) and (x_0, y_0) becomes larger.

Property 4: The asymmetry value for an SCTG distribution should generally be at its maximum for the correct value of θ_l in Eq. (10).

For an estimated optical center, the angle θ_l defines the line that divides the assignment of pixels to clockwise or counterclockwise SCTGs, as illustrated in Fig. 2(b). As shown in the analysis of Property 2, $\psi_t^{\mathcal{Y}}(r', \theta')$ is always non-negative given the true value of θ . An

incorrect value of θ_l , however, leads to an erroneous line and wrong assignments of pixels to clockwise or counterclockwise SCTGs. A wrongly assigned pixel has an opposite sign for its SCTG, and therefore it will be non-positive. These non-positive values decrease the positive skew of an SCTG distribution and reduce its asymmetry. Likewise, it can be seen that smaller errors in θ_l lead to fewer pixels with an incorrect SCTG sign, and thus greater asymmetry of the SCTG distribution. As a result, the asymmetry value is generally maximized with the correct value of θ_l .

4.4.2 Experimental Validation—We have compared the numerical asymmetry values of over 65 natural vignetted images before and after vignetting correction (the process of obtaining the ground truth will be explained in Sec. 6.1). For each of the images, three estimates of the optical center were randomly chosen with a distance of 0, 10, and 20 pixels from the true optical center. The asymmetry values of the SCTG distributions were computed with each of these three hypothesized optical centers. Before vignetting correction, significant increases in asymmetry values were observed for greater distances from the true optical center, from an average of 0.05 for a correctly estimated center, to 0.14 for a 10-pixel error and 0.19 for a 20-pixel error. By contrast, after the vignetting correction no obvious changes were observed (0.054 vs. 0.051 vs. 0.057). That a greater estimation error generally corresponds to a higher asymmetry value in a vignetted image is a favorable property for optimization of the optical center, which is described in Sec. 5.1.

5 Estimation

In this section, we apply the distribution properties of the RG and SCTG described in Sec. 4 to estimate the optical center and vignetting effect given a single image. To obtain reliable estimates of the optical center, we propose a nonlinear optimization based algorithm that minimizes the asymmetry of the SCTG distribution. With the estimated optical center, we present two variants of vignetting estimation that minimize the asymmetry of the RG distribution: (1) a least-squares solution that determines the vignetting value at discrete radii, and (2) a nonlinear optimization solution for fitting a parametric model of vignetting.

5.1 Optical Center Estimation with SCTG

As described in Sec. 4.4, the SCTG distribution is more symmetric with a more accurate estimate of θ_l and optical center. Therefore, our method solves for the optical center (x_0, \hat{y}_0) via the following minimization:

$$(\hat{x}_0, \hat{y}_0) = \arg \min_{(x'_0, y'_0)} \arg \max_{\theta_l} \Gamma_{st}(Z), \quad (21)$$

where the distribution asymmetry Γ_{st} of the SCTG is defined in Eq. (15).

The minimization in Eq. (21) can be accomplished by repeating two processes of estimating θ_l by fixing (x'_0, y'_0) and estimating (x'_0, y'_0) by fixing θ_l , until convergence. We carry out the former process with the scheme for determining θ_l explained in Sec. 3.3. For the latter process, we utilize the Levenberg-Marquardt (L-M) algorithm [26] with (x'_0, y'_0) initialized as the numerical center of the image coordinates. As shown in Fig. 10, the profile of the

asymmetry measure with respect to the unknowns x'_0 and y'_0 enables quick convergence by L-M optimization to an accurate solution.

5.2 Vignetting Estimation with RG

We propose two variants for single image vignetting correction 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 nonlinear optimization.

5.2.1 Least-squares Solution with Discrete Radii—Our goal is to find the optimal vignetting function V that minimizes asymmetry of the RG distribution. As in Sec. 4.4.1, we denote $\ln Z$, $\ln I$, and $\ln V$ by \mathcal{Z} , \mathcal{I} , and \mathcal{V} , respectively. We also 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. By taking the logarithm of Eq. (16), we have

$$\psi_r^{\mathcal{Z}}(r, \theta) = \psi_r^{\mathcal{I}}(r, \theta) + \psi_r^{\mathcal{V}}(r, \theta). \quad (22)$$

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

$$\mathcal{V} = \arg \max_{\mathcal{V}} P(\mathcal{V} | \mathcal{Z}) \propto \arg \max_{\mathcal{V}} P(\mathcal{Z} | \mathcal{V}) P(\mathcal{V}). \quad (23)$$

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} :

$$P(\mathcal{V}) = e^{-\lambda_s \sum_{r_t \in S_r} \mathcal{V}''(r_t)^2}, \quad (24)$$

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 Eq. (22) we have

$$\psi_r^{\mathcal{Z}}(r, \theta) = \psi_r^{\mathcal{I}}(r, \theta) - \psi_r^{\mathcal{V}}(r, \theta). \quad (25)$$

We impose the commonly-used gradient sparsity prior [16], [15] on the vignetting-free image \mathcal{I} :

$$P(\mathcal{Z}|\mathcal{V}) = P(\psi_r^{\mathcal{Z}}) = e^{-|\psi_r^{\mathcal{Z}}|^\alpha}, \quad \alpha < 1. \quad (26)$$

Here, we use $|\psi_r^{\mathcal{Z}}|$ because of symmetry of the RG distribution for \mathcal{Z} .

Substituting Eq. (25) into Eq. (26), we have

$$P(\mathcal{Z}|\mathcal{V}) = e^{-\sum_{(r,\theta)} |\psi_r^{\mathcal{Z}}(r,\theta) - \psi_r^{\mathcal{V}}(r)|^\alpha}, \quad (27)$$

where $\psi_r^{\mathcal{V}}(r) = (\mathcal{V}(r_m) - \mathcal{V}(r_{m-1})) / \delta r$ 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

$$\mathcal{O} = \sum_{(r,\theta)} |\psi_r^{\mathcal{Z}}(t, \theta) - \psi_r^{\mathcal{V}}(r)|^\alpha + \lambda_s \sum_{r_t \in S_r} \mathcal{V}''(r_t)^2 \quad (28)$$

where λ_s is empirically set to 0.1. Our goal here is to find the values of $\mathcal{V}(r_t)$, $t = \{0, 1, \dots, n-1\}$, that minimize \mathcal{O} .

To effectively apply this energy function, a proper sparsity parameter α for the RG distribution of \mathcal{Z} must be selected. As given in Eq. (26), α must be less than 1. However, very small values of α allow noise to more strongly bias the solution [45], [16]. 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, Eq. (28) does not have a closed-form solution.

To optimize Eq. (28), we employ an iteratively re-weighted least squares (IRLS) technique [15], [23]. 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

$$\mathcal{O}_k = \sum_{(r,\theta)} w_k(r, \theta) \left(\psi_r^{\mathcal{Z}}(r, \theta) - \psi_r^{\mathcal{V}_k}(r) \right)^2 + \lambda_s \sum_{r_t \in S_r} \mathcal{V}_k''(r_t)^2. \quad (29)$$

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

$$\begin{aligned} w_k(r, \theta) &= e^{-S_1} \left(1 - e^{-S_2} \right), \\ S_1 &= \left| \psi_r^{\mathcal{Z}}(r, \theta) - \psi_r^{\mathcal{V}_{k-1}}(r) \right|, \\ S_2 &= \alpha (S_1)^{\alpha-1}. \end{aligned} \quad (30)$$

Minimization of \mathcal{O}_k in Eq. (29) amounts to making the gradient of \mathcal{V} equal to the gradient of \mathcal{Z} for each pixel while enforcing local smoothness on \mathcal{V} . Directly solving Eq. (29) has no unique solution because adding a constant value to \mathcal{V} will not influence the value of \mathcal{O}_k . To

address this issue, we introduce another constraint as shown in the following new objective function:

$$\mathcal{O}_k = \sum_{(r,\theta)} w_k(r,\theta) \left(\psi_r^{\mathcal{L}}(r,\theta) - \psi_r^{\mathcal{V}_k}(r) \right)^2 + \lambda_s \sum_{r_t \in S_r} \mathcal{V}_k''(r_t)^2 + \epsilon \sum_{r_t \in S_r} \mathcal{V}_k(r_t)^2 \quad (31)$$

where ϵ is a very small value (e.g. 0.00001). Compared with [44], we add this constraint to make \mathcal{V} as small as possible (i.e., V in Eq. (16) as close to 1 as possible), leading \mathcal{V} to a unique solution. Without this constraint, the corresponding objective function in [44] is incapable of distinguishing between solutions which are different with a constant shift over the vignetting attenuation values.

The minimization of Eq. (31) relative to \mathcal{V} is a standard regularized least-squares problem. It can be accomplished by the ridge regression (or Tikhonov regularization) technique [3] and always has a closed-form solution.

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{L} 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 [15], the recomputed weights in our problem are always within the range $[0, 1]$. Fig. 11 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 (in Eq. (16)) and the radial gradient prior on \mathcal{L} . 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.

5.2.2 Model-based Solution—Many vignetting models exist, including polynomial functions [2], [28], hyperbolic cosine functions [38], 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], [10]. In this paper, we use the extended Kang-Weiss model [43] 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

$$\begin{aligned}
V(r) &= A(r)G(r), \quad (r) \in \Omega \\
A(r) &= \frac{1}{(1+(r/f)^2)^2}, \\
G(r) &= (1 - \alpha_1 r - \dots - \alpha_p r^p),
\end{aligned} \tag{32}$$

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

$$\mathcal{O} = \lambda \Gamma_r \left(\frac{Z}{V} \right) + (1 - \lambda) \left(\frac{N_b}{N_\Omega} \right)^{1/4}, \tag{33}$$

where $\Gamma_r \left(\frac{Z}{V} \right)$ is the measure of asymmetry for image Z/V using Eq. (15), 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 Eq. (33) penalizes outlier pixels.

To find the optimal vignetting model, we minimize the energy function in Eq. (33) using the Levenberg-Marquardt (L-M) algorithm [26]. We first solve for the focal length by fixing the geometric factor G to be 0. We then fix the 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 Eq. (32). 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. 12, the profile of the energy function (Eq. (33)) 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.

6 Results

We applied our algorithms on images captured using a Canon G3, Canon EOS 20D, Nikon E775, HP 945, as well as on images from the Berkeley Segmentation Dataset [21] and from flickr® (<http://www.flickr.com/>). To obtain a linear camera response, the single-image radiometric calibration method in [17] was applied as a preprocessing step to our algorithms.

6.1 Ground Truth Measurements

To obtain ground truth measurements of the optical center for each camera and setting, we use the DLR Camera Calibration Toolbox [29]. Ten images of a chessboard-like calibration panel were captured from different views and distances at each camera setting. The landmarks/corners of the panel were then detected with the DLR CalDe tool and some

manual interaction. From the image coordinates of these calibration features, the intrinsic and extrinsic parameters of the camera were estimated with the DLR CalLab tool.

The ground truth vignetting was obtained using the approach described in [43], [42]. Specifically, for each camera setting, a series of images of a distant white surface under approximately uniform illumination was captured, joined and blended to produce an accurate calibration image. The ratio of intensity with respect to the optical center in this image reflects the vignetting effect.

6.2 Estimation of Optical Center

We ran our algorithm for optical center estimation on a variety of simulated and real images.

Simulated Data—We first simulated different vignetting effects using the Kang-Weiss model (as described in Sec. 5.2.2) by keeping its off-axis illumination factor and neglecting its geometric factor. We set $f = \{250, 500, 1300, 2000, 3000\}$ (in pixels) and added the simulated vignetting effects with optical centers randomized on circles centered at the numerical center of the image coordinates. These circles were sampled at four different radii (5, 15, 25 and 30 pixels), giving a total of 20 simulated vignetting effects. These effects were applied to a set of 65 real-world images considered to be vignetting-free, as they were captured with a large focal length. The simulated vignetting effects were added to each of the vignetting-free images by multiplying the original image intensities by the vignetting attenuation value. We show the estimation errors of our approach in Fig. 13, where the error value for each combination of f and optical center shift is averaged over the 65 images.

Since our algorithm relies on vignetting information to estimate the optical center, it becomes less effective for images which are weakly vignettted, as shown in Fig. 13. On the other hand, the estimation is shown to be fairly insensitive to different shifts in the optical center away from the numerical center of the images.

Real images—We also applied our algorithm to real images taken by the aforementioned cameras. The focal length and zoom were varied to obtain different optical centers [37], and we additionally varied the focus and aperture size to produce different vignetting characteristics. For each camera setting, we captured 15 images of real-world scenes. We ran our algorithm on each of these images, and present comparisons to ground truth values.

We found that the zoom parameter has little influence on accuracy. In contrast, the focal length, focus distance, and aperture setting have more obvious effects, as shown in Tab. 1 based on the 15 images. We only report the results for the Canon EOS 20D in Tab. 1. Results for the other cameras have very similar error statistics. Camera settings that generate greater vignetting (shorter focal length, larger focus distance, and larger aperture size) usually lead to more accurate estimates of the optical center. With our technique, good accuracy can be obtained even for complex scenes.

6.3 Estimation of Vignetting

Indoor images—We ran our vignetting correction algorithms on 20 indoor images. The vignetting artifacts in indoor images are generally difficult to correct due to greater

illumination non-uniformity [43]. Since our methods are based on modeling the asymmetry of the gradient distributions instead of the intensity distributions, they work well for indoor images. The results shown in the top rows of Fig. 14 demonstrate that our methods are able to effectively reduce vignetting despite highly non-uniform illumination.

Texture images—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 distribution 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. 14(a).

Images from Berkeley Segmentation Dataset and flickr®—We ran our algorithms on the Berkeley Segmentation Dataset and on 57 images from flickr® that exhibit obvious vignetting. Vignetting was observed to be significantly reduced for photographs with very different content. Results for the most heavily vignettted images in the Berkeley Dataset are shown in Fig. 15. Our techniques were also found to be effective for flickr® images (Fig. 16), as indicated by the flatter brightness profiles of the vignetting corrected images in comparison to the original images. The lower flickr® image in Fig. 16 demonstrates the usefulness of our techniques to correct vignetting also in microscopy imaging. Vignetting correction provides an effective means of optimising digital microscopy image appearances for printing, image analysis, or telepathology [14].

Quantitative Evaluation and Comparison—To quantitatively evaluate the accuracy of our algorithms and compare them with previous methods, we ran our algorithms and Zheng *et al.*'s algorithm [43] on 80 images captured from outdoor (36), indoor (29) and texture (15) scenes, respectively. For all of the images, the ground truth vignetting is known. For outdoor scenes, our model-fitting variant performs the best while the method of Zheng *et al.* and our least-squares variant are comparable in performance.

We found that for indoor and texture scenes, our two methods, in 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.* [43] depends 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 nonuniform illumination while the method of Zheng *et al.* is less able to, as shown in Fig. 17. We observed that the algorithm of Zheng *et al.* [43] may generate segments that are inhomogeneous or non-uniformly textured, which leads to sub-optimal results.

Benefits of Optical Center Estimation—To validate the benefits of optical center estimation on vignetting correction, we ran our least squares vignetting correction algorithm on 20 real-world images captured from each of the four test cameras, with imaging parameters set to generate significant vignetting. We found that the mean/standard-deviation

of the Mean Squared Errors ($\times 10^{-3}$) of the vignetting estimations decrease from 2.08/1.72 (with the optical center assumed to be at the image center) to 1.29/0.64 (with the optical center estimated by the proposed algorithm).

6.4 Speed

We compared the speed of our vignetting methods to that of Zheng *et al.*'s algorithm [43] on a total of 80 outdoor, indoor, and textured images. All images were resized to a resolution of 450×600 (which does not affect the special shape of image gradient distributions as described in Sec. 2.3), and all algorithms were implemented in Matlab (except for the segmentation component of [43] in C++) and run on a Dell PC with a 2.39 GHz Intel Core 2 CPU. A substantial increase in speed is likely to be possible with an optimized C++ implementation.

Our algorithms achieved on average a speed-up of 3-5 times compared with Zheng *et al.*'s algorithm, from an average execution time of 285 seconds for Zheng *et al.*'s algorithm to 64 seconds of our least squares technique and 80 seconds for our model-based algorithm. Our methods are faster mainly because they do not require an iterative segmentation process. Note that the timings of our methods include the average execution time (29 sec) spent on automatic estimation of the optical center.

7 Discussion

Our proposed vignetting correction techniques work well for both non-texture images and texture images. There are, however, failure cases, due to our assumption of symmetric distribution of edges. This means that our techniques are less effective on images with predominantly inhomogeneous textures, e.g., a close view of trees. Failure may also result in cases where the dominant brightness variation is caused by outdoor illumination (e.g., bright sun in the sky) or where there is inhomogeneous shading/reflectance in indoor images. Highly noisy images (e.g., under low-light conditions) will also not be properly corrected due to the possible skewed gradient distribution. In such cases, the algorithm of Zheng *et al.* [43], [42] is expected to be more robust due to the use of a robust statistical function.

Our model-based vignetting variant uses a small number of parameters, and as such, has a better chance of converging to an optimal solution. However, since its optimization is nonlinear, convergence is slower than the least squares variant. Also, not all images with vignetting fit the Kang-Weiss vignetting model. Cameras with specially designed lenses, for example, may produce vignetting effects that deviate from this model. Here, the more flexible least squares variant would perform better.

8 Conclusion

We presented novel techniques for single-image vignetting correction based on the symmetric distribution of two particular image gradients: the semicircular tangential gradient (SCTG) and radial gradient (RG). Our techniques model the symmetry of the RG and SCTG distributions over the entire image and avoid the segmentation that is required by previous single-image methods. Moreover, our techniques are able to achieve higher

accuracies by determining the optical center precisely from a single input image. Experiments on a wide range of natural images show that our techniques are overall more robust and accurate, particularly for images with textures and non-uniform illumination, which are difficult to handle effectively using segmentation-based approaches. Our methods are also faster than the previous segmentation-based approach, achieving a speed-up of 3-5 times, and with comparable or better results.

We would like to extend the techniques in this paper to more general global variations of image brightness, e.g., nonuniformity artifacts in medical imaging [40], [39], to the simultaneous estimation of optical center, vignetting and other camera parameters (e.g., the camera response curve [7]), and to other applications such as high dynamic range imaging [25], radiometric calibration of agricultural cameras [24], and medical imaging [40], [39].

Finally, the implementation of our algorithms and some examples presented in this paper are available for public usage at <https://www.eecis.udel.edu/wiki/vims/index.php/Main/Software>.


Acknowledgments

This work was made possible by support from the National Science Foundation (NSF) grants CDI Type I GEO1124664, ANT0636726 and from the National Institute of Health (NIH) via grants EB006266, DA022807, NS045839, and NS065347.

Biography



Yuanjie Zheng is a senior research investigator in Perelman School of Medicine at the University of Pennsylvania and the primary contact of Image Analysis Core at the Penn Vision Research Center. He received his Ph.D. degree in Pattern Recognition and Intelligent Systems from Shanghai Jiao Tong University of China in 2006. His current research interests are in computer vision, medical image analysis and clinical study.



Stephen Lin received the BSE degree in electrical engineering from Princeton University, New Jersey, and the PhD degree in computer science and engineering from the University of Michigan, Ann Arbor. He is a senior researcher in the Internet Graphics group, Microsoft

Research Asia. He served as a Program Co-Chair for the International Conference on Computer Vision 2011 and the Pacific-Rim Symposium on Image and Video Technology 2009. His research interests include computer vision, image processing, and computer graphics. He is a member of the IEEE.



Sing Bing Kang is a principal researcher at Microsoft Research. He received his Ph.D. in robotics from Carnegie Mellon University, Pittsburgh in 1994. His areas of interest are computer vision and computer graphics, more specifically image-based modeling as well as image and video enhancement. Sing Bing has co-edited two books ("Panoramic Vision" and "Emerging Topics in Computer Vision") and co-authored two books ("Image-Based Rendering" and "Image-Based Modeling of Plants and Trees"). He has served as area chair and member of technical committee for the major computer vision conferences (ICCV, CVPR, ECCV). In addition, he has served as papers committee member for SIGGRAPH and SIGGRAPH Asia. Sing Bing was program chair for ACCV 2007 and CVPR 2009, and is currently Associate Editor-in-Chief for IEEE Transactions on Pattern Recognition and Machine Intelligence. He is an IEEE Fellow.



Rui Xiao received the PhD degree in Biostatistics from the University of Michigan, Ann Arbor in 2010. She is currently an assistant professor in the Department of Biostatistics and Epidemiology at the University of Pennsylvania School of Medicine. Her primary research interests are in the design, implementation, and analysis of neuroimaging related studies and high throughput genetic studies.



James Gee, Ph.D., is Associate Professor of Radiologic Science and Computer and Information Science, Director of the the HHMI-NIBIB (Howard Hughes Medical Institute National Institutes of Biomedical Imaging and Bioengineering) Interfaces Program in Biomedical Imaging and Informational Sciences, and Co-Director of the Translational Biomedical Imaging Center of the Institute for Translational Medicine and Therapeutics, all at the University of Pennsylvania, Philadelphia. Dr. Gee's group, the Penn Image Computing and Science Laboratory (PICSL), conducts basic and applied research in mathematical, computational and statistical methods for extracting, analyzing and integrating information from imaging and ancillary data. Of particular interest to the lab are methods for detecting, quantifying and modeling the ways in which anatomy and physiology can vary in nature, over time, or as a consequence of disease or therapy. A major goal of this research is its translation into practical tools, made freely and publicly available through the group's open-source ITK-SNAP, ANTS, DTI-TK and ITK software, which have consistently ranked as the best performing and most widely used applications in segmentation, registration, diffusion tensor imaging and morphometry. The group's portfolio of interdisciplinary collaborations spans different model systems and the major modalities in biological and medical imaging, in integrative studies of structure-function relationships of the brain, lung, heart and musculoskeletal system in health and disease.



Chandra Kambhamettu received his PhD in Computer Science and Engineering from the University of South Florida in 1994. From 1994-96, he was a research scientist at NASA-Goddard, where he received the “1995 Excellence in Research Award.” He is currently a Professor in the Department of Computer Science, University of Delaware, where he leads the Video/Image Modeling and Synthesis (VIMS) group. His research interests include video modeling and image analysis for biomedical, remote sensing and multimedia applications. Dr. Kambhamettu is best known for his work in motion analysis of deformable bodies, for which he received NSF CAREER award in 2000. He published over 170 peer-reviewed papers, supervised ten PhD students and several Masters students in his areas of interest. He served as Area Chair, and has been technical committee member for leading computer vision and medical conferences. He also served as associate editor for Pattern Recognition journal and IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI), and currently serving as associate editor for Pattern Recognition Letters journal.

REFERENCES

1. Apostoloff N, Fitzgibbon A. Bayesian video matting using learnt image priors. CVPR. 2004

2. Asada N, Amano A, Baba M. Photometric calibration of zoom lens systems. Proc. Int. Conf. on Pattern Recognition. 1996:186–190.
3. Draper, NR.; Smith, H. Applied Regression Analysis. 2nd edition. John Wiley; New York: 1981.
4. Faugeras OD, Luong Q. t. Camera self-calibration: theory and experiments. ECCV. 1992:321–334.
5. Fergus R, Singh B, Hertzmann A, Roweis ST, Freeman WT. Removing camera shake from a single photograph. ACM Transactions on Graphics. 2006; 25(3):787–794.
6. Field DJ. What is the goal of sensory coding? Neural Computation. 1994; 6(4):559–601.
7. Goldman DB. Vignette and exposure calibration and compensation. IEEE Transactions on Pattern Analysis and Machine Intelligence. 2010; 32:2276–2288. [PubMed: 20975123]
8. Hartley RI, Hayman E, Agapito LD, Reid I. Camera calibration and the search for infinity. ICCV. 1999:510–517.
9. Juang R, Majumder A. Photometric self-calibration of a projector-camera system. CVPR. 2007
10. Kang S, Weiss R. Can we calibrate a camera using an image of a flat textureless lambertian surface? European Conf. on Computer Vision. 2000; II:640–653.
11. Kim SJ, Pollefeys M. Robust radiometric calibration and vignetting correction. IEEE Transactions on Pattern Analysis and Machine Intelligence. 2008; 30(4):562–576. [PubMed: 18276964]
12. Kuthirummal S, Agarwala A, Goldman DB, Nayar SK. Priors for large photo collections and what they reveal about cameras. ECCV. 2008
13. Lenz RK, Tsai RY. Techniques for calibration of the scale factor and image center for high accuracy 3-d machine vision metrology. IEEE Transactions on Pattern Analysis and Machine Intelligence. 1988; 10(15):713–720.
14. Leong FJW-M, Brady M, McGee JO. Correction of uneven illumination (vignetting) in digital microscopy images. Journal of Clinical Pathology. 2003; 56:619–621. [PubMed: 12890815]
15. Levin A, Fergus R, Durand F, Freeman WT. Image and depth from a conventional camera with a coded aperture. ACM Trans. on Graphics. 2007; 26(3)
16. Levin A, Zomet A, Weiss Y. Learning to perceive transparency from the statistics of natural scenes. NIPS. 2002
17. Lin S, Gu J, Yamazaki S, Shum HY. Radiometric calibration using a single image. CVPR. 2004; 2:938–945.
18. Lin, Y-H.; Low, TW. Device and method for optical center detection.. 2007. US Patent Patent No.: US7,307,709 B2
19. Litvinov A, Schechner Y. Addressing radiometric nonidealities: A unified framework. CVPR. 2005:52–59.
20. Lyu S. Single image vignetting correction with natural image statistics in derivative domains. Proc. ICIP'10. 2010
21. Martin D, Fowlkes C, Tal D, Malik J. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. ICCV. Jul. 2001 2:416–423.
22. Matsushita Y, Lin S. Radiometric calibration from noise distributions. CVPR. 2007
23. Meer, P. Robust techniques for computer vision. Prentice-Hall; 2005. p. 107-190.
24. Olsen D, Dou C, Zhang X, Hu L, Kim H, Hildum E. Radiometric calibration for agcam. Remote Sensing. 2010; 2:464–477.
25. Pouli T, Cunningham DW, Reinhard E. A survey of image statistics relevant to computer graphics. Computer Graphics Forum. 2011; 00:1–28.
26. Press, WH.; Teukolsky, SA.; Vetterling, WT.; Flannery, BP. Numerical Recipes in C: The Art of Scientific Computing. Cambridge University Press; New York, NY, USA: 1992.
27. Roth S, Black M. Fields of experts: A framework for learning image priors. CVPR. 2005:860–867.
28. Sawchuk AA. Real-time correction of intensity nonlinearities in imaging systems. IEEE Trans. on Computers. 1977; 26(1):34–39.
29. Strobl, KH.; Sepp, W.; Fuchs, S.; Paredes, C.; Arbter, K. DLR CalDe and DLR CalLab. Institute of Robotics and Mechatronics, German Aerospace Center (DLR);

30. Tappen M, Russell B, Freeman W. Exploiting the sparse derivative prior for super-resolution and image demosaicing. *IEEE Workshop on Statistical and Computational Theories of Vision*. 2003
31. Torralba A, Oliva A. Statistics of natural image categories. *Network: Computation in Neural System*. 2003; 14:391–412.
32. Tsai RY. A versatile camera calibration technique for high-accuracy 3d machine vision metrology using off-the-shelf tv cameras and lenses. *IEEE Journal of Robotics and Automation*. 1987; 3(4): 323–344.
33. Uyttendaele M, Criminisi A, Kang SB, Winder S, Szeliski R, Hartley R. Image-based interactive exploration of real-world environments. *IEEE Computer Graphics and Applications*. 2004; 24:52–63. [PubMed: 15628073]
34. Wang L-L, Tsai W-H. Computing camera parameters using vanishing-line information from a rectangular parallelepiped. *Machine Vision and Applications*. 1990; 3(3):129–141.
35. Weiss Y. Deriving intrinsic images from image sequences. *ICCV*. 2001
36. Willson, RG. PhD thesis. Carnegie Mellon University; Pittsburgh, PA.: 1994. Modelling and calibration of automated zoom lenses..
37. Willson RG, Shafer SA. What is the center of the image? *Journal of the Optical Society of America A*. 1994; 11(11):2946–2955.
38. Yu W. Practical anti-vignetting methods for digital cameras. *IEEE Trans. on Cons. Elect*. 2004; 50:975–983.
39. Zheng Y, Gee J. Automatic correction of intensity nonuniformity from sparseness of gradient distribution in medical images. *CVPR*. 2010
40. Zheng Y, Grossman M, Awate S, Gee J. Automatic correction of intensity nonuniformity from sparseness of gradient distribution in medical images. *the 12th International Conference on Medical Image Computing and Computer Assisted Intervention*. 2009
41. Zheng Y, Kambhamettu C, Lin S. Single-image optical center estimation from vignetting and tangential gradient symmetry. *CVPR*. 2009
42. Zheng Y, Lin S, Kambhamettu C, Kang SB. Single-image vignetting correction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2009; 31(12):2243–2256. [PubMed: 19834144]
43. Zheng Y, Lin S, Kang SB. Single-image vignetting correction. *CVPR*. 2006
44. Zheng Y, Yu J, Kang SB, Lin S, Kambhamettu C. Single-image vignetting correction using radial gradient symmetry. *CVPR*. 2008
45. Zhu SC, Mumford D. Prior learning and gibbs reaction-diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 1997; 19(11):1236–1250.




Fig. 1.

Radial gradient ψ_r^Z is defined as the projection of the conventional image gradient ∇Z on the radial direction r .




Fig. 2.

Illustration of TG and SCTG. (a) TG Ψ_t^Z is defined as the projection of the conventional image gradient ∇Z on the tangential direction t . (b) The line through the true optical center (green dot) (x_0, y_0) and its estimation (pink dot) (x'_0, y'_0) divides the blue circle into two parts. With Eq. (10), the clockwise SCTG is assigned to pixels on the part with the green arrow while the counterclockwise SCTG is assigned to the part with the blue arrow. (c) SCTG assignments are different for different optical center estimates.

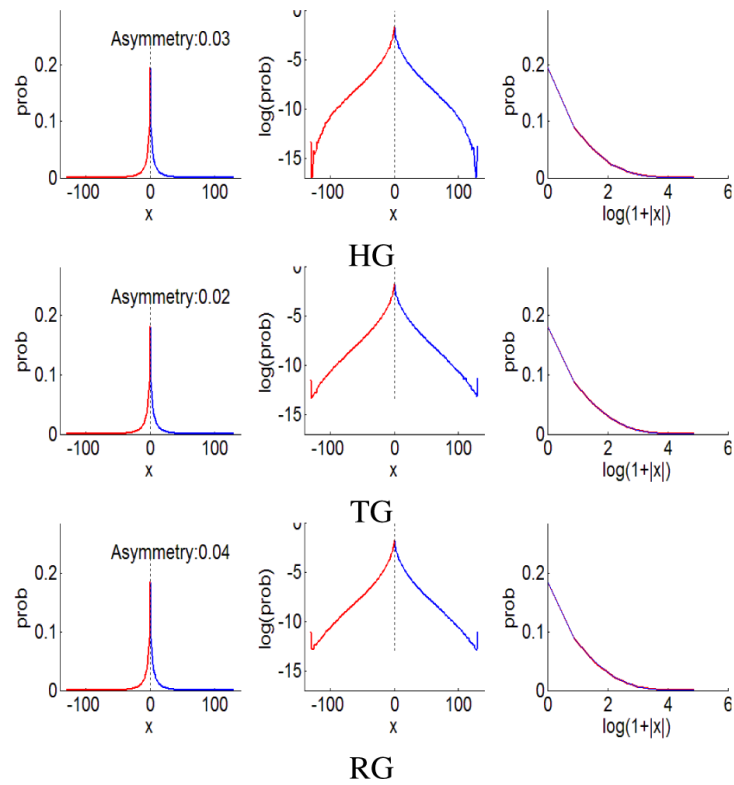


Fig. 3. Statistical regularities of an ensemble of natural scene images free from vignetting or other directional brightness variation, demonstrated with HG, TG and RG histograms. From left to right: regular histogram, log-histogram, and $\log(1 + |x|)$ -histogram.

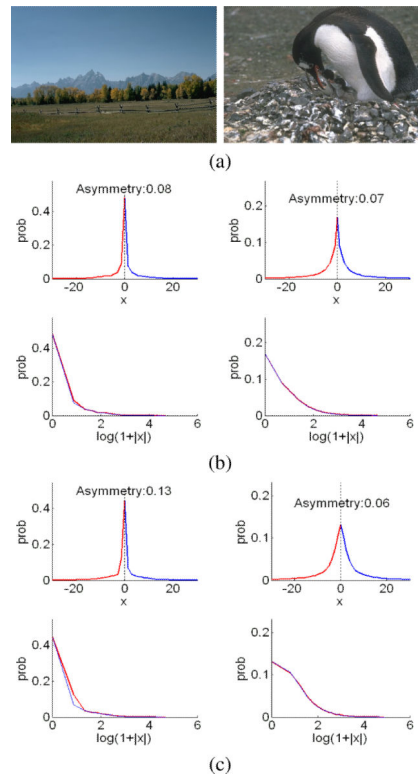


Fig. 4. Gradient histograms for two natural images in (a). In (b) and (c), top to bottom: regular histogram and corresponding $\log(1 + |x|)$ histogram. (b) displays plots for HG while (c) shows plots for RG.

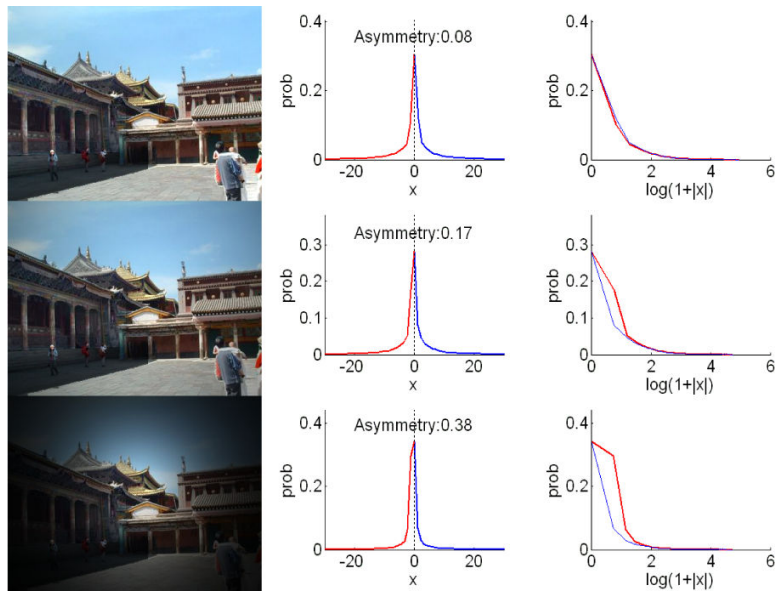


Fig. 5. Comparison of asymmetry of RG distribution for varying degrees of vignetting. From left to right: image, regular histogram with asymmetry value, and $\log(1 + |x|)$ histogram. From top to bottom: increasing degrees of vignetting.

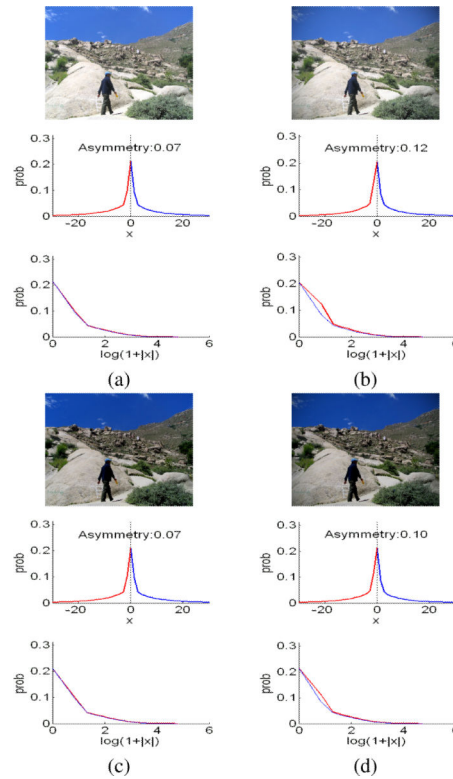


Fig. 6. Effect of darker images on asymmetry. (a) Original image, (b) image with synthetic vignetting, (c) darkened version of (a), (d) same amount of synthetic vignetting applied to (c). For each of (a)-(d), from top to bottom: image, regular histogram with asymmetry value, $\log(1 + |x|)$ histogram. Notice that brighter images with vignetting have a greater asymmetry value.

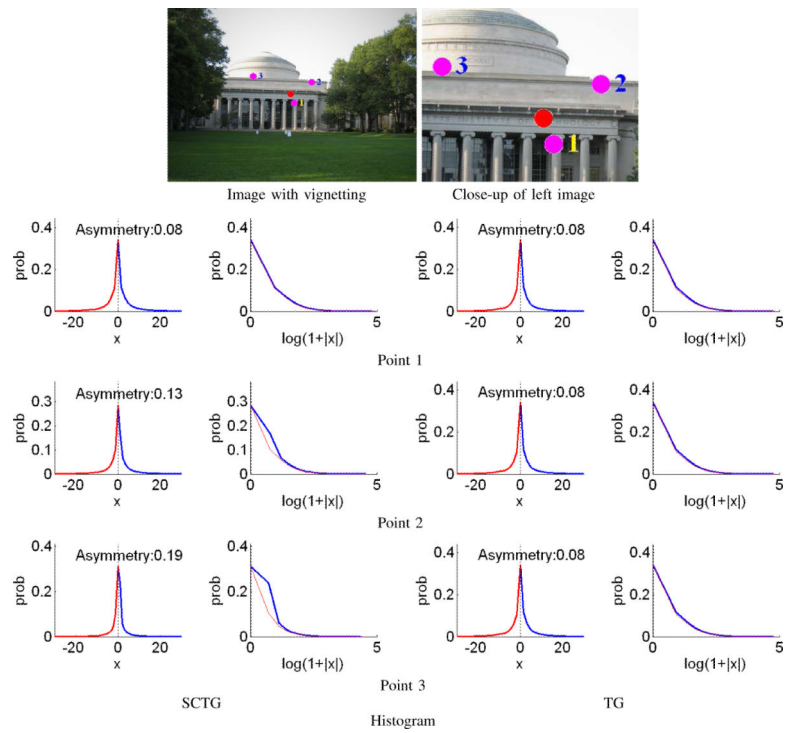


Fig. 7. Comparison of SCTG and TG distributions. The true optical center and three estimated optical centers are marked by the red dot and three purple dots, respectively. For each estimate, we show the regular histogram with asymmetry value and the $\log(1 + |x|)$ histogram.

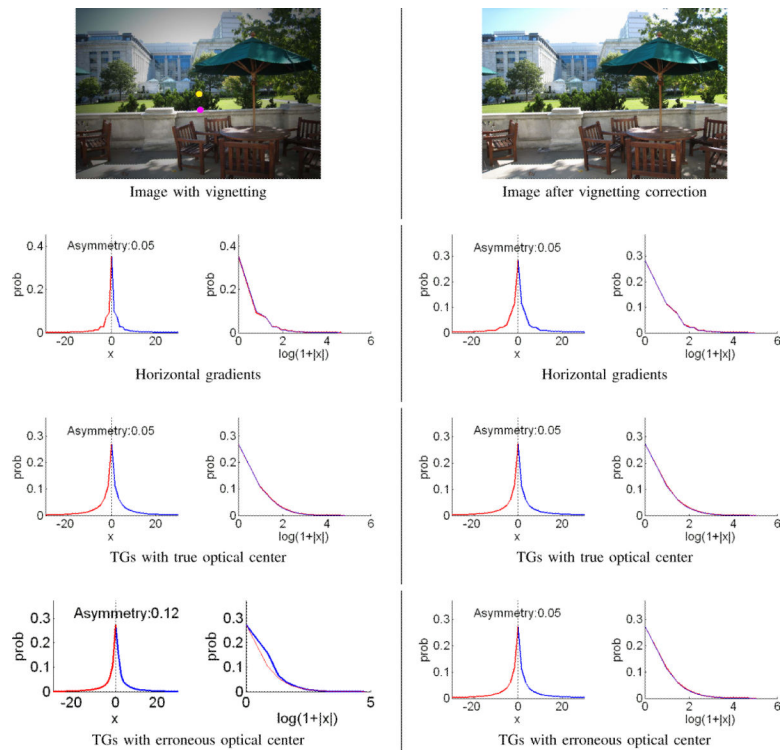


Fig. 8. Comparison of SCTG and HG distributions. In each image, the true optical center and estimated optical center are marked by the yellow dot and purple dot, respectively. With each distribution, a $\log(1 + |x|)$ histogram is also shown to emphasize asymmetry.

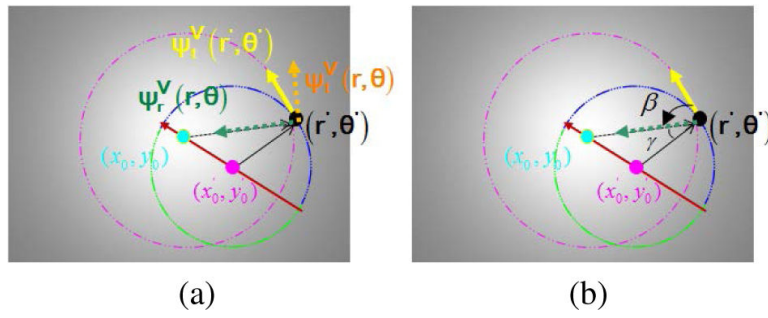


Fig. 9.
Geometric analysis of SCTG asymmetry caused by vignetting.




Fig. 10. Profile of asymmetry measure with respect to different hypothesized positions of the optical center (x_0, y_0) for the image in Fig. 7. The true optical center is marked by the black dot.

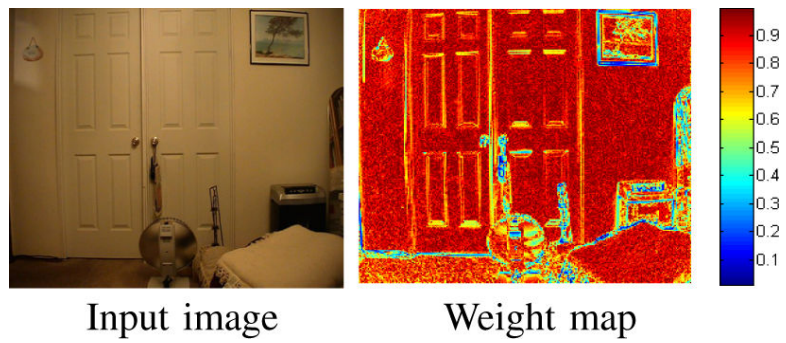


Fig. 11. Computed weights in the least squares variant (Eq. (29)) after the 3rd iteration of the IRLS algorithm.

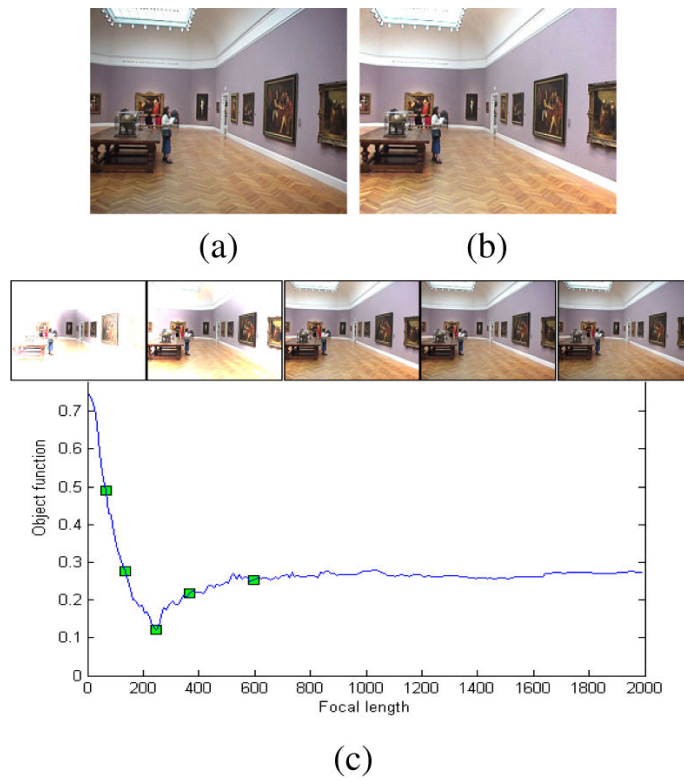


Fig. 12. Model-based vignetting correction. (a) Input image, (b) final corrected image, and (c) graph of objective function (Eq. (33)) vs. focal length. The images above the graph, from left to right, correspond to corrected versions using focal length values indicated by green squares on the curve. The focal length yielding the minimum value is the final solution.


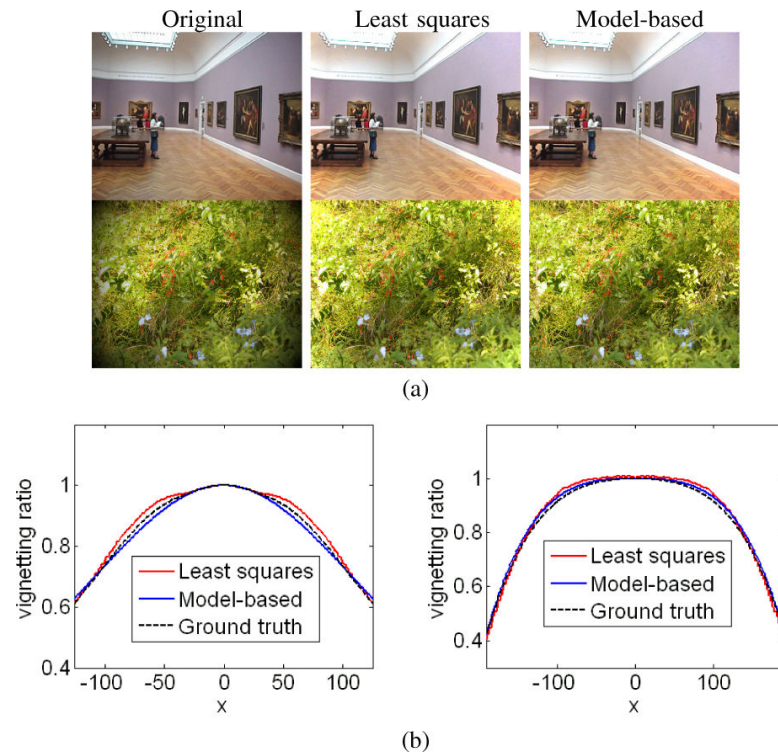


Fig. 13. Errors in optical center estimation using images with simulated vignetting. The vignetting effects are simulated with different focal lengths f and random shifts of the optical center away from the numerical center of the image coordinates.

**Fig. 14.**

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 the upper and lower images in (a), respectively. The red and blue curves are obtained by our least squares and model-based methods, respectively, and the black dotted curves are the ground truth.




Fig. 15. The four most heavily vignettted images in the Berkeley Segmentation Dataset [21] and the vignetting correction results with our techniques.

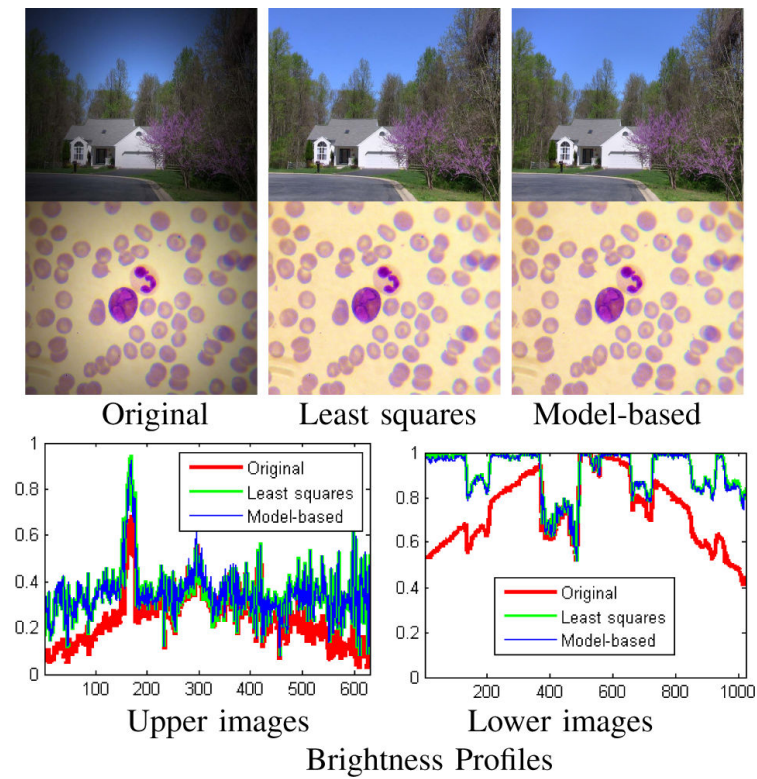


Fig. 16. Vignetting correction results using our methods on two images taken from flickr®. Brightness profiles are drawn for pixels on a middle horizontal line, from left to right. The lower original image is from microscopy imaging.

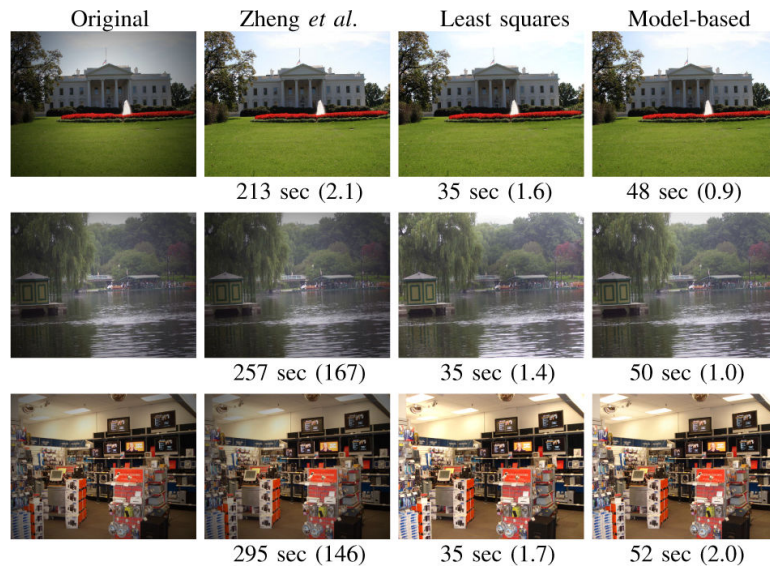


Fig. 17. Comparisons of speed and accuracy. The numbers in parentheses are mean squared errors ($\times 10^{-3}$).

TABLE 1

Error statistics (mean μ and standard deviation σ) with a Canon EOS 20D at different zooms z (with focal length 50mm), focal lengths f (with 27mm-136mm zoom), focus distances d (with 100mm zoom), and aperture sizes a (with 100mm zoom).

Zoom	z	27mm	40mm	90mm	136mm
Error	μ	1.0	1.3	1.1	1.0
	σ	0.6	1.9	0.7	0.7
Focal length	f	33mm	50mm	100mm	135mm
Error	μ	1.4	1.8	4.7	6.3
	σ	0.8	1.0	2.1	3.3
Focus	d	1m	5m	10m	∞
Error	μ	14.3	9.6	5.7	1.2
	σ	31.3	5.5	3.1	0.8
Aperture	a	f 5.6	f 4.5	f 3.5	f 2.8
Error	μ	14.1	13.3	8.4	1.5
	σ	21.0	18.9	7.8	0.9

Errors are measured in pixels.

TABLE 2

Comparison of mean/standard-deviation of the Mean Squared Errors ($\times 10^{-3}$) for 80 images.

	Zheng <i>et al.</i>	Least squares	Model-based
Outdoor	1.9/0.7	1.8/0.9	1.3/0.3
Indoor	3.1/1.9	2.4/1.5	2.6/1.3
Texture	5.9/2.3	5.2/2.1	3.9/1.6