

# Single-Image Vignetting Correction

Yuanjie Zheng\*  
Shanghai Jiaotong University

Stephen Lin  
Microsoft Research Asia

Sing Bing Kang  
Microsoft Research

## Abstract

*In this paper, we propose a method for determining the vignetting function given only a single image. Our method is designed to handle both textured and untextured regions in order to maximize the use of available information. To extract vignetting information from an image, we present adaptations of segmentation techniques that locate image regions with reliable data for vignetting estimation. Within each image region, our method capitalizes on frequency characteristics and physical properties of vignetting to distinguish it from other sources of intensity variation. The vignetting data acquired from regions are weighted according to a presented reliability measure to promote robustness in estimation. Comprehensive experiments demonstrate the effectiveness of this technique on a broad range of images.*

## 1. Introduction

Vignetting refers to the phenomenon of brightness attenuation away from the image center, and is an artifact that is prevalent in photography. Although not objectionable to the average viewer at low levels, it can significantly impair computer vision algorithms that rely on precise intensity data to analyze a scene. Applications in which vignetting distortions can be particularly damaging include photometric methods such as shape from shading, appearance-based techniques such as object recognition, and image mosaicing.

Several mechanisms may be responsible for vignetting effects. Some arise from the optical properties of camera lenses, the most prominent of which is off-axis illumination falloff or the  $\cos^4$  law. This contribution to vignetting results from foreshortening of the lens when viewed from increasing angles from the optical axis [7]. Other sources of vignetting are geometric in nature. For example, light arriving at oblique angles to the optical axis may be partially obstructed by the field stop or lens rim.

To determine the vignetting effects in an image, the most straightforward approach involves capturing an image completely spanned by a uniform scene region, such that brightness variations can solely be attributed to vignetting [12, 1, 6, 14]. In such a calibration image, ratios of intensity with respect to the pixel on the optical axis describe the vignetting function. Suitable imaging conditions for this approach, however, can be challenging to produce due to uneven illumination and camera tilt, and the vignetting measurements are valid only for images captured by the camera under the same camera settings. Moreover, a calibration image can be recorded only if the camera is at hand; consequently, this approach cannot be used to correct images captured by unknown cameras, such as images downloaded from the web.

A vignetting function can alternatively be computed from image sequences with overlapping views of an arbitrary static scene [5, 9, 4]. In this approach, point correspondences are first determined in the overlapping image regions. Since a given scene point has a different position in each image, its brightness may be differently attenuated by vignetting. From the aggregate attenuation information from all correspondences, the vignetting function can be accurately recovered without assumptions on the scene.

These previous approaches require either a collection of overlapping images or an image of a calibration scene. However, often in practice only a single image of an arbitrary scene is available. The previous techniques gain information for vignetting correction from pixels with equal scene radiance but differing attenuations of brightness. For a single arbitrary input image, this information becomes challenging to obtain, since it is difficult to identify pixels having the same scene radiance while differing appreciably in vignetting attenuation.

In this paper, we show that it is possible to correct or reduce vignetting given just a single image. To maximize the use of available information in the image, our technique extracts vignetting information from *both textured and untextured regions*. Large image regions appropriate for vignetting function estimation are identified by proposed adaptations to segmentation methods. To counter the adverse effects of vignetting on segmentation, our method

---

\*This work was done while Yuanjie Zheng was a visiting student at Microsoft Research Asia.




Figure 1. Tilt angles  $\tau$  and  $\chi$  in the Kang-Weiss vignetting model.

iteratively re-segments the image with respect to progressively refined estimates of the vignetting function. Additionally, spatial variations in segmentation scale are used in a manner that enhances collection of reliable vignetting data.

In extracting vignetting information from a given region, we take advantage of physical vignetting characteristics to diminish the influence of textures and other sources of intensity variation. With the joint information of disparate image regions, we describe a method for computing the vignetting function. The effectiveness of this vignetting correction method is supported by experiments on a wide variety of images.

## 2. Vignetting model

Most methods for vignetting correction use a parametric vignetting model to simplify estimation and minimize the influence of image noise. Typically used are empirical models such as polynomial functions [4, 12] and hyperbolic cosine functions [14]. Models based on physical considerations include that of Asada *et al.* [1], which accounts for off-axis illumination and light path obstruction, and that of Kang and Weiss [6] which additionally incorporates scene-based tilt effects. Tilt describes intensity variations within a scene region that are caused by differences in distance from the camera, i.e., closer points appear brighter due to the inverse square law of illumination. Although not intrinsic to the imaging system, the intensity attenuation effects caused by tilt must be accounted for in single-image vignetting estimation.

Besides having physically meaningful parameters, an important property of physical models is that their highly structured and constrained form facilitates estimation in cases where data is sparse and/or noisy. In this work, we use an extension of the Kang-Weiss model, originally designed for a single planar surface of constant albedo, to multiple surfaces of possibly different color. Additionally, we generalize its linear model of geometric vignetting to a polynomial form.

### 2.1. Kang-Weiss model

We consider an image with zero skew, an aspect ratio of 1, and principal point at the image center with image coordinates  $(u, v) = (0, 0)$ . In the Kang-Weiss vignetting model [6], brightness ratios are described in terms of an off-axis illumination factor  $A$ , a geometric factor  $G$ , and a tilt factor  $T$ . For a pixel  $i$  at  $(u_i, v_i)$  with distance  $r_i$  from the image center, the vignetting function  $\varphi$  is expressed as

$$\varphi_i = A_i G_i T_i = \vartheta_{r_i} T_i \quad \text{for } i = 1 \cdots N, \quad (1)$$

where

$$\begin{aligned} A_i &= \frac{1}{(1 + (r_i/f)^2)^2}, \\ G_i &= (1 - \alpha_1 r_i), \\ \vartheta_{r_i} &= A_i G_i, \\ T_i &= \cos \tau \left( 1 + \frac{\tan \tau}{f} (u_i \sin \chi - v_i \cos \chi) \right)^3. \end{aligned} \quad (2)$$

$N$  is the number of pixels in the image,  $f$  is the effective focal length of the camera, and  $\alpha_1$  represents a coefficient in the geometric vignetting factor. The tilt parameters  $\chi, \tau$  respectively describe the rotation angle of a planar scene surface around an axis parallel to the optical axis, and the rotation angle around the  $x$ -axis of this rotated plane, as illustrated in Fig. 1.

The model  $\varphi$  in Eq. (1) can be decomposed into the global vignetting function  $\vartheta$  of the camera and the natural attenuation caused by local tilt effects  $T$  in the scene. Note that  $\vartheta$  is rotationally symmetric; thus, it can be specified as a 1D function of the radial distance  $r_i$  from the image center.

### 2.2. Extended vignetting model

In an arbitrary input image, numerous regions with different local tilt factors may exist. To account for multiple surfaces in an image, we present an extension of the Kang-Weiss model in which different image regions can have different tilt angles. The tilt factor of Eq. (2) is modified to

$$T_i = \cos \tau_{s_i} \left( 1 + \frac{\tan \tau_{s_i}}{f} (u_i \sin \chi_{s_i} - v_i \cos \chi_{s_i}) \right)^3, \quad (3)$$

where  $s_i$  indexes the region containing pixel  $i$ .

We also extend the linear geometric factor to a more general polynomial form:

$$G_i = (1 - \alpha_1 r_i - \cdots - \alpha_p r_i^p), \quad (4)$$

where  $p$  represents a polynomial order that can be arbitrarily set according to a desired precision. This generalized representation provides a closer fit to the geometric vignetting effects that we have observed in practice. In contrast to using a polynomial as the overall vignetting model, representing only the geometric component by a polynomial allows the overall model to explicitly account for local tilt effects and global off-axis illumination.

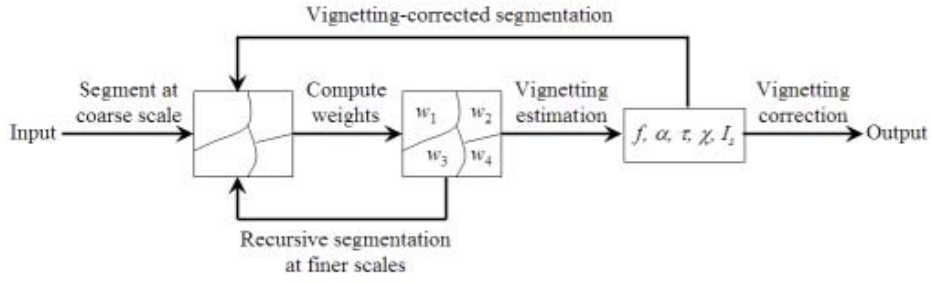


Figure 3. Overview of vignetting function estimation.

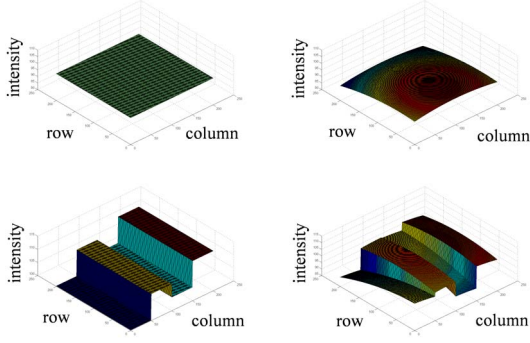


Figure 2. Vignetting over multiple regions. Top row: without and with vignetting for a single uniform region. Bottom row: without and with vignetting for multiple regions.

### 2.3. Vignetting energy function

Let the scene radiance  $I_s$  of a region  $s$  be expressed by its ratio  $\lambda_s$  to the scene radiance  $I_0$  of the center pixel, i.e.,  $I_s = \lambda_s I_0$ . Given an image with  $M$  regions of different scene radiance, we formulate the vignetting solution as the minimization of the following energy function:

$$E = \sum_{s=1}^M \sum_{i=1}^{N_s} w_i (\lambda_s I_0 T_i \vartheta_{r_i} - z_i)^2, \quad (5)$$

where  $i$  indexes the  $N_s$  pixels in region  $s$ ,  $z_i$  is the pixel value in the vignettted image, and  $w_i$  is a weight assigned to pixel  $i$ . In color images,  $z$  represents an RGB vector. For ease of explanation, we express  $z$  in this paper as a single color channel, and overall energies are averaged from separate color components.

In this energy function, the parameters to be estimated are the focal length  $f$  in the off-axis component, the  $\alpha$  coefficients of the geometric factor, the tilt angles  $\tau_s$  and  $\chi_s$ , the scene radiance of the center pixel  $I_0$ , and the radiance ratio  $\lambda_s$  of each region. In processing multiple image regions as illustrated in Fig. 2, minimization of this energy function can intuitively be viewed as simultaneously solving for local region parameters  $I_s$ ,  $\tau_s$  and  $\chi_s$  that give a smooth alignment of vignetting attenuations between regions, while

optimizing the underlying global vignetting parameters  $f$ ,  $\alpha_1, \dots, \alpha_p$ .

With the estimated parameters, the vignetting corrected image is then given by  $z_i / \vartheta_{r_i}$ . We note that the estimated local tilt factors may contain other effects that can appear similar to tilt, such as non-uniform illumination or shading. *In the vignetting corrected image, these tilt and tilt-like factors are all retained so as not to produce an unnatural-looking result.* Only the attenuation attributed to the imaging system itself (off-axis illumination and geometric factors) is corrected.

In this formulation, the scene is assumed to contain some piecewise planar Lambertian surfaces that are uniformly illuminated and occupy significant portions of the image. Although typical scenes are considerably more complex than uniform planar surfaces, we will later describe how vignetting data in an image can be separated from other intensity variations such as texture, and how the weights  $w$  are set to enable robust use of this energy function.

### 3. Algorithm overview

The high-level flow of our algorithm is illustrated in Fig. 3. In each iteration through the procedure, the image is first segmented at a coarse scale, and for each region a reliability measure of the region data for vignetting estimation is computed. For regions that exhibit greater consistency with physical vignetting characteristics and with other regions, a higher reliability weight is assigned. Low weights may indicate regions with multiple distinct surfaces, so these regions are recursively segmented at incrementally finer scales until weights of the smaller regions exceed a threshold or regions becomes negligible in size. With this segmentation approach, the segmentation scale varies spatially in a manner that facilitates collection of vignetting data.

After spatially adaptive segmentation, regions with high reliability weights are used to estimate the vignetting model parameters. Since the preceding segmentations may be corrupted by the presence of vignetting, the subsequent iteration of the procedure re-computes segmentation boundaries from an image corrected using the currently estimated vignetting model. Better segmentation results lead to im-

proved vignetting estimates, and these iterations are repeated until the estimates converge.

The major components of this algorithm are described in the following sections.

## 4. Vignetting-based image segmentation

To obtain information for vignetting estimation, pixels having the same scene radiance need to be identified in the input image. Our method addresses this problem with proposed adaptations to existing segmentation methods. To facilitate the location of reliable vignetting data, segmentation scales are spatially varied over the image, and the adverse effects of vignetting on segmentation are progressively reduced as the vignetting function estimate is refined.

### 4.1. Spatial variations in scale

Sets of pixels with the same scene radiance provide more valuable information if they span a broader range of vignetting attenuations. In the context of segmentation, larger regions are therefore preferable. While relatively large regions can be obtained with a coarse segmentation scale, many of these regions may be unreliable for vignetting estimation since they may contain multiple surfaces or include areas with non-uniform illumination. In an effort to gain useful data from an unreliable region, our method recursively segments it into smaller regions that potentially consist of better data for vignetting estimation. This recursive segmentation proceeds until regions have a high reliability weight or become of negligible size according to a threshold of 225 pixels used in our implementation. Regions of very small size generally contain insignificant changes in vignetting attenuation, and the inclusion of such regions would bias the optimization process.

In the recursive segmentation procedure, incrementally finer scales of segmentation are used. For methods such as mean shift [3] and region competition [13], segmentation scale is essentially controlled by a parameter on variation within each feature class, where a feature may simply be pixel intensity or color. With such approaches, a finer partitioning of a low-weight region can be obtained by segmenting the region with a decreased parameter value. In other techniques such as graph cuts [15] and Blobworld [2], the degree of segmentation is set according to a given number of feature classes in an image. There exist various ways to set the number of classes, including user specification, data clustering, and minimum description length criteria [11]. For recursive segmentation, since each region belongs to a certain class, a finer partitioning of the region can be obtained by segmenting it with the number of feature classes specified as two.

With this general adaptation, segmentation scale varies over an image in a manner designed to maximize the qual-

ity of vignetting data. In our implementation, we employ graph cut segmentation [15] with per-pixel feature vectors composed of six color/texture attributes. The color components are the RGB values, and the local texture descriptors are the polarity, anisotropy and normalized texture contrast described in [2].

### 4.2. Accounting for vignetting

Two pixels of the same scene radiance may exhibit significantly different image intensities due to variations in vignetting attenuation. In segmentation, a consequence of this vignetting is that a homogeneous scene area may be divided into separate image regions. Vignetting may also result in heterogeneous image areas being segmented together due to lower contrasts at greater radial distances. For better stability in vignetting estimation, the effects of vignetting on segmentation should be minimized.

To address vignetting effects in segmentation, after each iteration through the procedure in Fig. 3, the estimated vignetting function is accounted for in segmentations during the subsequent iteration. Specifically, the vignetting corrected image computed with the currently estimated parameters is used in place of the original input image in determining segmentation boundaries. The corrected image is used only for segmentation purposes, and the colors in the original image are still used for vignetting estimation.

As the segmentations improve from reduced vignetting effects, the estimated vignetting function also is progressively refined. This process is repeated until the difference between vignetting functions in consecutive iterations falls below a prescribed threshold, where the difference is measured as

$$\Delta\vartheta = \frac{1}{k} \sum_r \|\vartheta_r(t) - \vartheta_r(t-1)\|. \quad (6)$$

$\vartheta(t)$  represents the global vignetting function at iteration  $t$ , and radial distances  $r$  are sampled at  $k$  uniform intervals, where  $k = 100$  in our implementation.

## 5. Region weighting

To guide the vignetting-based segmentation process and promote robust vignetting estimation, the reliability of data in each image region is evaluated and used as a region weight. A region is considered to be reliable if it exhibits consistency with physical vignetting characteristics and conforms to vignetting observed elsewhere in the image.

Initially, no vignetting estimates are known, so reliability is measured in the first iteration of the algorithm according to how closely the region data can be represented by our physically-based vignetting model. For a given region, an estimate  $\vartheta'$  of the vignetting function is computed similarly

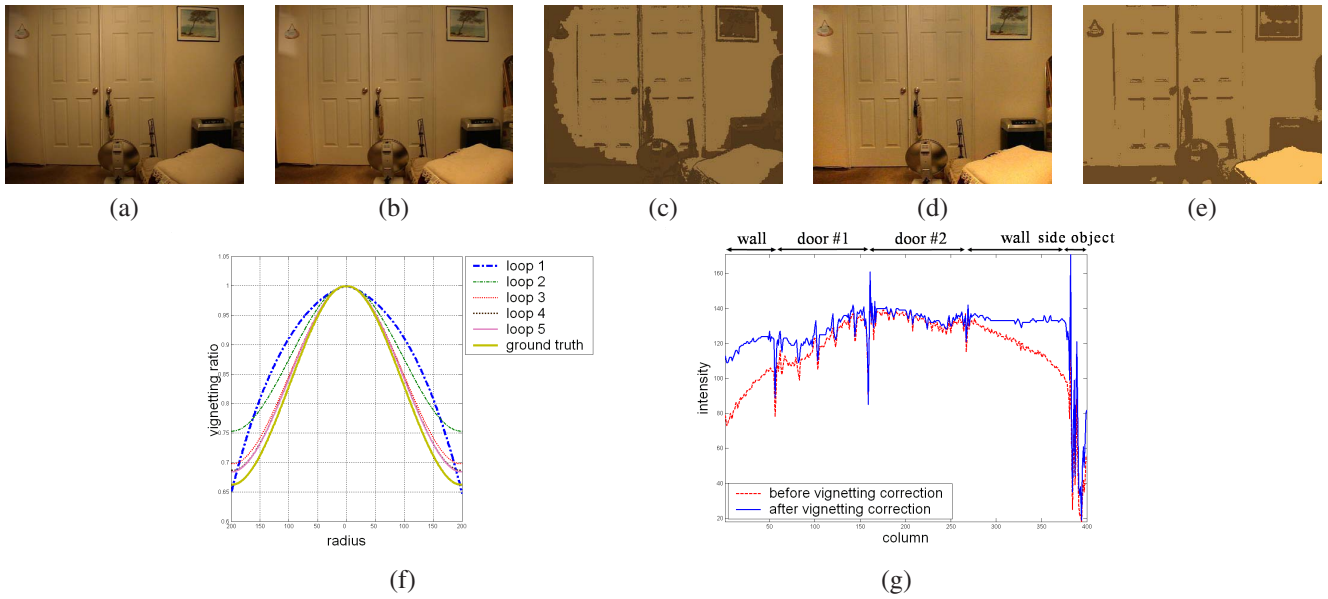


Figure 4. Effects of vignetting compensation in segmentation. (a) Original image; (b) Vignetting correction with segmentation that does not account for vignetting; (c) Segmentation without accounting for vignetting; (d) Vignetting correction with segmentation that accounts for vignetting; (e) Segmentation that accounts for vignetting; (f) Estimated vignetting functions after each iteration in comparison to the ground truth; (g) Intensity profile before (red) and after (blue) correction, shown for the image row that passes through the image center.

to the technique described in Section 6, and the weight for region  $s$  is computed as

$$w_s = \frac{1}{N_s} \exp \left\{ - \sum_{i=1}^{N_s} \left\| \vartheta'_{r_i} - \frac{z_i}{\lambda_s I_0 T_i} \right\| \right\}. \quad (7)$$

Each pixel is assigned the weight of its region.

The presence of texture in a region does not preclude it from having a high weight. In contrast to textures which typically exhibit high frequency variations, vignetting is a low frequency phenomenon with a wavelength on the order of the image width. This difference in frequency characteristics allows vignetting effects to be discerned in many textured regions.

At the end of each iteration, an estimate of the vignetting function is determined and used as  $\vartheta'$  in the following iteration. As the vignetting parameters are progressively refined, computed weights will more closely reflect the quality of region data. In cases where the texture or shading in a region coincidentally approximates the characteristics of vignetting, it will be assigned a low weight if it is inconsistent with the vignetting observed in other parts of the image.

## 6. Vignetting estimation

For a collection of segmented regions, the many unknown parameters create a complicated solution space. To simplify optimization, we use a stepwise method for parameter initialization prior to estimating the vignetting function.

In the first step, initial values of relative scene radiances  $\lambda_s$  are determined for each region without consideration of vignetting and tilt parameters. For pixels  $i$  and  $j$  at the same radius  $r$  but from different regions, their vignetting attenuation should be equal, so their image values  $z_i$  and  $z_j$  should differ only in scene radiance. Based on this property, relative scene radiance values are initialized by minimizing the function

$$E_1 = \sum_r \sum_{r_i, r_j=r; s_i \neq s_j} w_i w_j \left( \frac{z_i}{\lambda_{s_i}} - \frac{z_j}{\lambda_{s_j}} \right)^2.$$

The  $\lambda_s$  values are solved in the least squares sense by singular value decomposition (SVD) on a system of equations  $\sqrt{w_i w_j} \left( \frac{z_i}{\lambda_{s_i}} - \frac{z_j}{\lambda_{s_j}} \right) = 0$  where  $\frac{1}{\lambda_{s_i}}$  and  $\frac{1}{\lambda_{s_j}}$  are unknowns. To expedite minimization of this function, a set of pixels at a given radius and within the same region may be represented by a single pixel with the average color of the set.

With the initial values of  $\lambda_s$ , the second step initializes the parameters  $f$ ,  $I_0$ , and  $\alpha_1, \dots, \alpha_p$ , where  $p$  is the polynomial order used in the geometric factor of Eq. 4. Ignoring local tilt factors, this is computed with the energy function

$$E_2 = \sum_{s=1}^M \sum_{i=1}^{N_s} w_i (\lambda_s I_0 \vartheta_{r_i} - z_i)^2. \quad (8)$$

This function is iteratively solved by incrementally increasing the polynomial order from  $k = 1$  to  $k = p$ , and using the previously computed polynomial coefficients  $\alpha_1, \dots, \alpha_{k-1}$

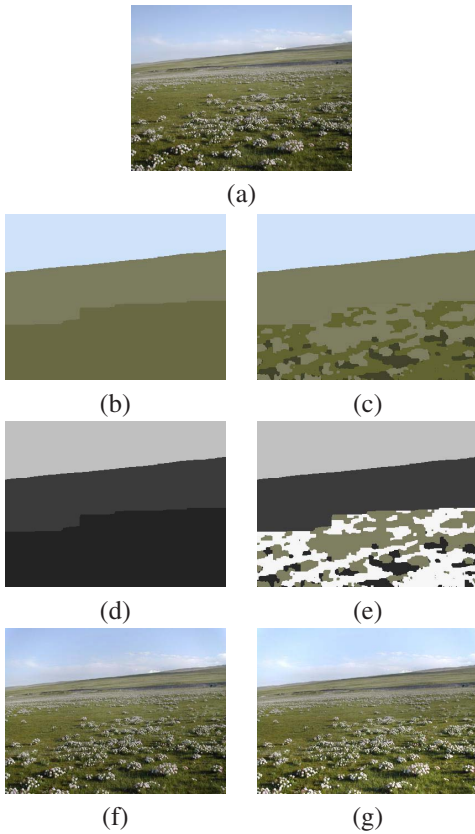


Figure 5. Recursive segmentation on low-weight regions. (a) Original image; (b) Regions prior to recursive segmentation; (c) Regions after recursive segmentation of one region; (d) Region weights of (b), where higher intensity indicates a higher weight; (e) Region weights of (c); (f) Vignetting correction result using region information of (b); (g) Vignetting correction result using region information of (c).

as initializations. In our implementation, we use a polynomial order of  $p = 4$ .

In the third step, the local tilt parameters  $\tau_s, \chi_s$  are estimated by optimizing the energy function in Eq. 5 with the other parameters fixed to their initialization values. After this initialization stage, all the parameters are jointly optimized in Eq. 5 to finally estimate the vignetting function. The optimizations of Eq. 5 and Eq. 8 are computed using the Levenberg-Marquardt algorithm [10].

## 7. Results

Our algorithm was evaluated on images captured with a Canon G3, Canon EOS 20D, and a Nikon E775. To obtain a linear camera response, the single-image radiometric calibration method of [8] can be applied as a preprocessing step to our algorithm. Ground truth vignetting functions of the cameras at different focal lengths were computed from multiple images of a distant white surface under approxi-

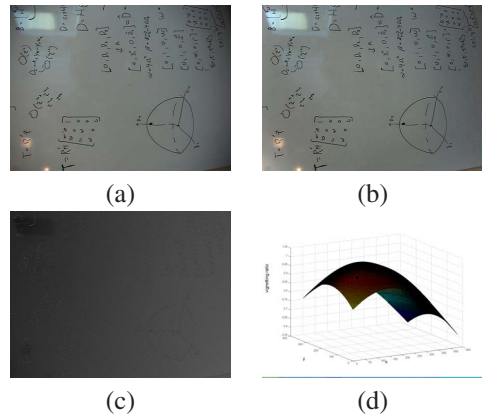


Figure 6. Tilt effects in vignetting estimation. (a) Original image with vignetting and tilt; (b) Image corrected for only vignetting using the proposed method; (c) Tilt image, where brighter areas indicate more distant points on a surface; (d) Estimated attenuation function with both vignetting and tilt.

Image	(a)	(b)	(c)	(d)	(e)	(f)
Error	1.373	2.988	1.812	2.327	2.592	0.973
Image	(g)	(h)	(i)	(j)	(k)	(l)
Error	2.368	3.176	0.823	2.782	2.501	1.473

Table 1. Error ( $\times 10^{-3}$ ) in estimated vignetting function for images in Fig. 7.

mately uniform illumination. A distant surface was used to minimize tilt effects, but generally a distant surface does not fully cover the image plane. We captured multiple images with camera translation such that each image pixel views the surface in at least one view. The image fragments of the white surface were joined and blended to obtain an accurate calibration image.

We first examine the effects of the proposed segmentation adaptations. Accounting for vignetting in segmentation leads to progressive improvements in the estimated vignetting function as exemplified in Fig. 4. The correction and segmentation results in (b) and (c) without vignetting compensation are equivalent to those after a single pass through the overall procedure shown in Fig. 3. With additional iterations, the enhanced segmentations lead to vignetting estimates that trend towards the ground truth.

The effect of recursive segmentation on a given region is illustrated in Fig. 5. Further segmentation of a low-weight region can produce sub-regions of higher weight. With this improvement in data quality, a more accurate vignetting function can be estimated.

While the goal of this work is to estimate and correct for the global vignetting function of the camera, tilt effects computed in vignetting estimation as shown in Fig. 6 could potentially provide some geometric information of

the scene. It should be noted though that estimated tilt values are only accurate for reliable regions with high weights.

Some vignetting correction results of our technique are presented in Fig. 7, along with the vignetting-based segmentation regions and their weights. Errors computed similarly to Eq. 6 between the estimated vignetting functions and the ground truth functions are listed in Table 1. While some slight vignetting artifacts may be visible under close examination, the correction quality is reasonable especially considering that only a single arbitrary input image is processed. For some indoor images, the amount of reliable data can possibly be low due to greater illumination non-uniformity. Images with poor data quality could potentially be identified within our method by examination of region weights, and indicated to the user.

In Fig. 8, we show the application of our method to image mosaicing. Even though vignetting correction was performed independently on each image of the sequence, a reasonable mosaicing result was still obtained. In cases where overlapping images are available, joint consideration of the vignetting data among all images in the sequence would likely lead to better results. In contrast to previous works on image mosaicing [5, 9, 4], our proposed method can also jointly process data from images containing *completely different content* if they are captured by the same camera under the same camera setting.

## 8. Conclusion

In this paper, we introduced a method for vignetting correction using only the information available in a single arbitrary image. Adaptations to general segmentation techniques are presented for locating regions with reliable vignetting data. Within an image region, the proposed method takes advantage of frequency characteristics and physical properties of vignetting to distinguish it from other sources of intensity variation. Experimental results demonstrate effective vignetting correction on a broad range of images.

Accurate correction results are generally obtained despite many regions having non-planar geometry and non-uniform illumination. The detrimental effects of non-planar geometry are reduced when distance variations of surface points from the camera are small in comparison to the distance of the surface itself, since variations in scene radiances become negligible. In many instances, the effects of non-uniform illumination appear similar to tilt, such that its effects on image intensity are incorporated into the estimated tilt factor. Low frequency vignetting effects also remain distinct when geometry and illumination exhibit texture-like high-frequency variations, such as among leaves on a tree. As a result, reliable vignetting data often exists even in image areas with significant geometry and illumination variation.

Directions for future work include joint estimation of



Figure 8. The image mosaic on top exhibits obvious vignetting effects. Below, the same sequence after vignetting is corrected separately in each image using our method. No image blending has been applied.

camera parameters such as principal point, aspect ratio, and skew, in addition to the vignetting function. In our current method, these camera parameters are assumed to be known from prior geometric calibration, but could potentially be recovered from vignetting information. Another interesting topic for future investigation is the examination of data in the RGB channels for region weighting, since vignetting should attenuate RGB values in a similar way, while other causes of region variation may not affect the channels equally.

## References

- [1] N. Asada, A. Amano, and M. Baba. Photometric calibration of zoom lens systems. In *ICPR 1996*, pages 186–190, 1996.
- [2] C. Carson, S. Belongie, H. Greenspan, and J. Malik. Blobworld: Image segmentation using expectation-maximization and its application to image querying. *IEEE Trans. PAMI*, 24(8):1026–1038, 2002.
- [3] D. Comaniciu and P. Meer. Mean shift: A robust approach toward feature space analysis. *IEEE Trans. PAMI*, 24(5):603–619, 2002.
- [4] D. B. Goldman and J. H. Chen. Vignette and exposure calibration and compensation. In *ICCV 2005*, pages 899–906, 2005.
- [5] J. Jia and C.-K. Tang. Tensor voting for image correction by global and local intensity alignment. *IEEE Trans. PAMI*, 27(1):36–50, 2005.
- [6] S. B. Kang and R. Weiss. Can we calibrate a camera using an image of a flat textureless lambertian surface? In *ECCV 2000*, volume II, pages 640–653, 2000.
- [7] M. V. Klein and T. E. Furtak. *Optics*. John Wiley and Sons, 1986.
- [8] S. Lin, J. Gu, S. Yamazaki, and H.-Y. Shum. Radiometric calibration using a single image. In *CVPR 2004*, volume 2, pages 938–945, 2004.

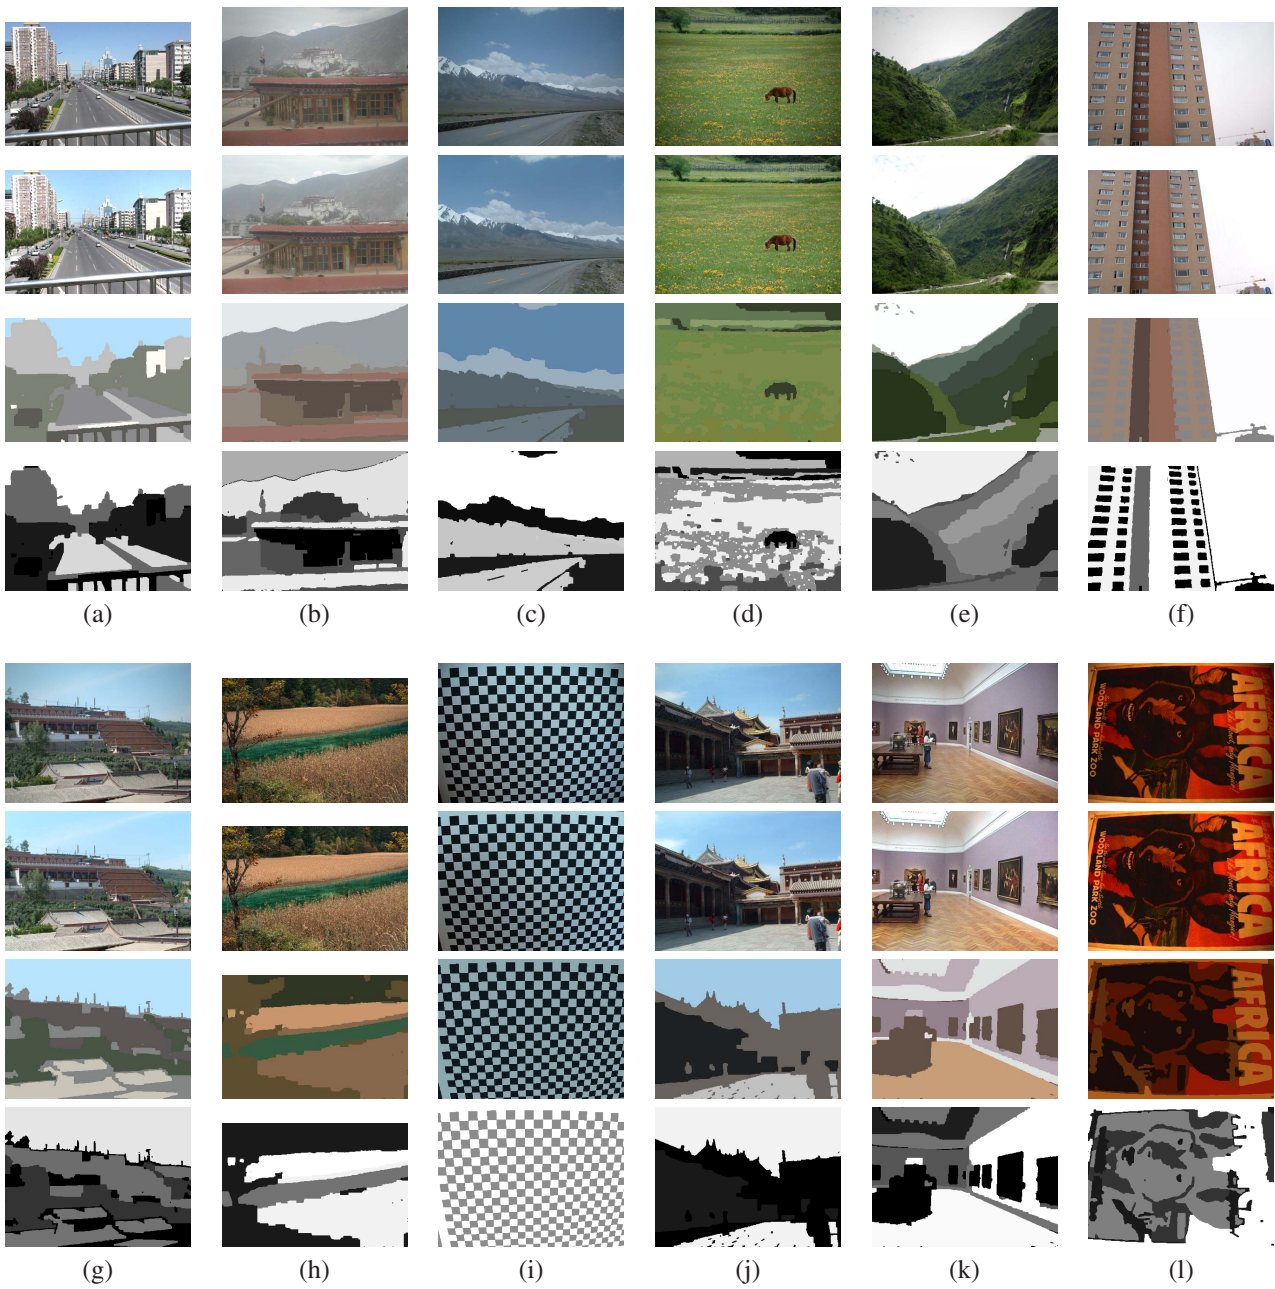


Figure 7. Vignetting correction results. Each set from top to bottom: original image, vignetting corrected image, vignetting-based segmentation, region weights (brighter pixels indicate higher weights).

[9] A. Litvinov and Y. Y. Schechner. Addressing radiometric nonidealities: A unified framework. In *CVPR 2005*, pages 52–59, 2005.

[10] 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, 1992.

[11] J. Rissanen. Modelling by shortest data description. *Automatica*, 14:465–471, 1978.

[12] A. A. Sawchuk. Real-time correction of intensity nonlinearities in imaging systems. *IEEE Trans. Computers*, 26(1):34–39, 1977.

[13] M. Tang and S. Ma. General scheme of region competition based on scale space. *IEEE Trans. PAMI*, 23(12):1366–1378, 2001.

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

[15] R. Zabih and V. Kolmogorov. Spatially coherent clustering using graph cuts. In *CVPR 2004*, volume 2, pages II–437–II–444, 2004.