Computational Aesthetics in Graphics, Visualization, and Imaging (2008)
P. Brown, D. W. Cunningham, V. Interrante, and J. McCormack (Editors)

Modeling and Synthesis of Aperture Effects in Cameras

Douglas Lanman[1,2], Ramesh Raskar[1,3], and Gabriel Taubin[2]

[1]Mitsubishi Electric Research Laboratories, Cambridge, MA, (USA)
[2]Brown University, Division of Engineering, Providence, RI (USA)
[3]Massachusetts Institute of Technology, Media Lab, Cambridge, MA (USA)

Abstract

In this paper we describe the capture, analysis, and synthesis of optical vignetting in conventional cameras. We analyze the spatially-varying point spread function (PSF) to accurately model the vignetting for any given focus or aperture setting. In contrast to existing "flat-field" calibration procedures, we propose a simple calibration pattern consisting of a two-dimensional array of point light sources – allowing simultaneous estimation of vignetting correction tables and spatially-varying blur kernels. We demonstrate the accuracy of our model by deblurring images with focus and aperture settings not sampled during calibration. We also introduce the Bokeh Brush*: a novel, post-capture method for full-resolution control of the shape of out-of-focus points. This effect is achieved by collecting a small set of images with varying basis aperture shapes. We demonstrate the effectiveness of this approach for a variety of scenes and aperture sets.*

Categories and Subject Descriptors (according to ACM CCS): I.3.8 [Computer Graphics]: Applications

1. Introduction

A professional photographer is faced with a seemingly great challenge: how to select the appropriate lens for a given situation at a moment's notice. While there are a variety of heuristics, such as the well-known "sunny f/16" rule, a photographer's skill in this task must be honed by experience. It is one of the goals of computational photography to reduce some of these concerns for both professional and amateur photographers. While previous works have examined methods for refocusing, deblurring, or augmenting conventional images, few have examined the topic of *bokeh*. In general, a good bokeh is characterized by a subtle blur for out-of-focus points – creating a pleasing separation between foreground and background objects in portrait or macro photography. In this paper we develop a new method to allow post-capture control of lens bokeh for still life scenes.

To inform our discussion of image bokeh, we present a unified approach to vignetting calibration in conventional cameras. Drawing upon recent work in computer vision and graphics, we propose a simple, yet accurate, vignetting and spatially-varying point spread function model. This model and calibration procedure should find broad applicability as more researchers begin exploring the topics of vignetting, highlight manipulation, and aesthetics.

1.1. Contributions

The vignetting and spatially-varying point spread function capture, analysis, and synthesis methods introduced in this paper integrate enhancements to a number of prior results in a novel way. The primary contributions include:

i. By exploiting the simple observation that the out-of-focus image of a point light directly gives the point spread function, we show a practical low-cost method to simultaneously estimate the vignetting and the spatially-varying point spread function. (Note that, while straightforward, this method can prove challenging in practice due to the long exposure times required with point sources.)
ii. We introduce the *Bokeh Brush*: a novel, post-capture method for full-resolution control of the shape of out-of-focus points. This effect is achieved by collecting a small set of images with varying basis aperture shapes. We demonstrate that optimal basis aperture selection is essentially a compression problem – one solution of which is to apply PCA or NMF to training aperture images. (Note that we assume a static scene so that multiple exposures can be obtained with varying aperture shapes.)

1.2. Related Work

The topic of vignetting correction can be subsumed within the larger field of radiometric calibration. As described in Litvinov and Schechner [LS05], cameras exhibit three primary types of radiometric non-idealities: (1) spatial non-uniformity due to vignetting, (2) nonlinear radiometric response of the sensor, and (3) temporal variations due to automatic gain control (AGC). Unfortunately, typical consumer-grade cameras do not allow users to precisely control intrinsic camera parameters and settings (e.g., zoom, focal length, and aperture). As a result, laboratory flat-field calibration using a uniform white light area source [Yu04] proves problematic – motivating recent efforts to develop simpler radiometric calibration procedures. Several authors have focused on single-image radiometric calibration, as well as single-image vignetting correction [ZLK06]. In most of these works the motivating application is creating image mosaics, whether using a sequence of still images [GC05, d'A07] or a video sequence [LS05].

Recently, several applications in computer vision and graphics have required high-accuracy estimates of spatially-varying point spread functions. Veeraraghavan et al. [VRA*07] and Levin et al. [LFDF07] considered *coded aperture imaging*. In those works, a spatially-modulated mask (i.e., an aperture pattern) was placed at the iris plane of a conventional camera. In the former work, a broadband mask enabled post-processing digital refocusing (at full sensor resolution) for layered Lambertian scenes. In the later work, the authors proposed a similar mask for simultaneously recovering scene depth and high-resolution images. In both cases, the authors proposed specific PSF calibration patterns, including: general scenes under natural image statistics and a planar pattern of random curves, respectively. We also recognize the closely-related work on confocal stereo and variable-aperture photography developed by Hasinoff and Kutulakos [HK06, HK07]. Note that we will discuss their models in more detail in Section 2.1.

2. Modeling Vignetting

Images produced by optical photography tend to exhibit a radial reduction in brightness that increases towards the image periphery. This reduction arises from a combination of factors, including: (1) limitations of the optical design of the camera, (2) the physical properties of light, and (3) particular characteristics of the imaging sensor. In this work we separate these effects using the taxonomy presented by Goldman and Chen [GC05] and Ray [Ray02].

Mechanical vignetting results in radial brightness attenuation due to physical obstructions in front of the lens body. Typical obstructions include lens hoods, filters, and secondary lenses. In contrast to other types of vignetting, mechanical vignetting can completely block light from reaching certain image regions, preventing those areas from being recovered by any correction algorithm.

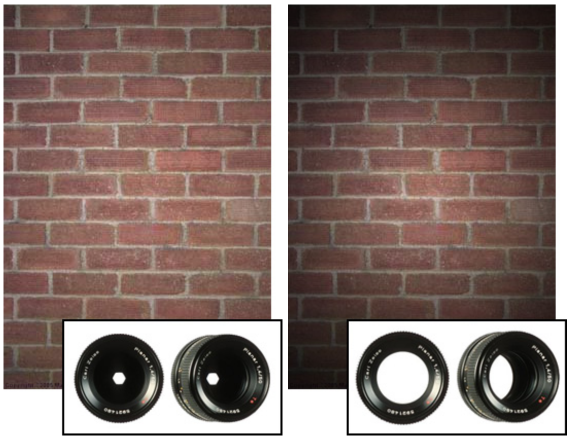

Figure 1: *Illustration of optical vignetting. From left to right: (a) reference image at f/5.6, (b) reference image at f/1.4, and (inset) illustration of entrance pupil shape as a function of incidence angle and aperture setting [vW07].*

Optical vignetting occurs in multi-element optical designs. As shown in Figure 1, for a given aperture the clear area will decrease for off-axis viewing angles and can be modeled using the variable cone model described in [AAB96]. Optical vignetting can be reduced by stopping down the lens (i.e., reducing the aperture size), since this will reduce exit pupil variation for large viewing angles.

Natural vignetting causes radial attenuation that, unlike the previous types, does not arise from occlusion of light. Instead, this source of vignetting is due to the physical properties of light and the geometric construction of typical cameras. Typically modeled using the approximate $\cos^4(\theta)$ law, where θ is the angle of light leaving the rear of the lens, natural vignetting combines the effects due to the inverse square fall-off of light, Lambert's law, and the foreshortening of the exit pupil for large incidence angles [KW00, vW07].

Pixel vignetting arises in digital cameras. Similar to mechanical and optical vignetting, pixel vignetting causes a radial falloff of light recorded by a digital sensor (e.g., CMOS) due to the finite depth of the photon well, causing light to be blocked from the detector regions for large incidence angles.

2.1. Geometric Model: Spatially-varying PSF

Recall that a thin lens can be characterized by

$$\frac{1}{f} = \frac{1}{f_D} + \frac{1}{D},$$

where f is the focal length, f_D is the separation between the image and lens planes, and D is the distance to the object plane (see Figure 2). As described by Bae and Durand [BD07], the diameter c of the PSF is given by

$$c = \frac{|S-D|}{S} \cdot \frac{f^2}{N(D-f)},$$

where S is the distance to a given out-of-focus point and N is the f-number. This model predicts that the PSF will scale as a function of the object distance S and the f-number N. As a result, a calibration procedure would need to sample both these parameters to fully characterize the point spread function. However, as noted by Hasinoff and Kutulakos [HK07], the effective blur diameter $\tilde{c}$ is given by the linear relation

$$\tilde{c} = \frac{|S-D|}{S}A,$$

where A is the aperture diameter. Under this approximation, we find that the spatially varying PSF could potentially be estimated from a single image. In conclusion, we find that the spatially-varying PSF $B(s,t;x,y)$ will scale linearly with the effective blur diameter $\tilde{c}$ such that

$$B_{\tilde{c}}(s,t;x,y) = \frac{1}{\tilde{c}^2} B_{\tilde{c}}\left(s,t;\frac{x}{\tilde{c}},\frac{y}{\tilde{c}}\right),$$

as given by Hasinoff and Kutulakos' model [HK07].

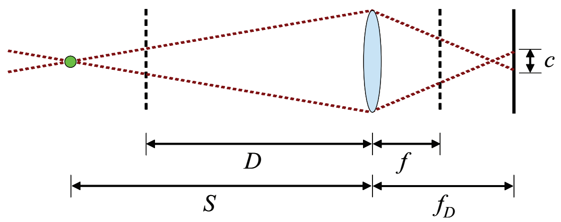

Figure 2: *The thin lens model. The aperture diameter is A and the focal length is f. The image plane and object plane distances are given by f_D and D, respectively. Out-of-focus points at S create a circle of confusion of diameter c [BD07].*

2.2. Photometric Model: Radial Intensity Fall-off

As shown in Figure 1, typical lenses demonstrate a significant radial fall-off in intensity for small f-numbers. While previous authors have fit a smooth function to a flat-field calibration data set [Yu04, AAB96], we propose a data-driven approach. For a small sampling of the camera settings, we collect a sparse set of vignetting coefficients in the image space. Afterwards, we apply scattered data interpolation (using radial basis functions) to determine the vignetting function for arbitrary camera settings and on a dense pixel-level grid (assuming the vignetting function is smoothly varying in both space and as a function of camera settings).

3. Data Capture

Given the geometric and photometric model in the previous section, we propose a robust method for estimating its parameters as a function of the general camera settings, including: zoom, focus, and aperture. In this paper, we restrict our analysis to fixed focal length lenses, such that the only intrinsic variables are: (1) the distance to the focus plane and (2) the f-number of the lens. In contrast to existing PSF and vignetting calibration approaches that utilize complicated area

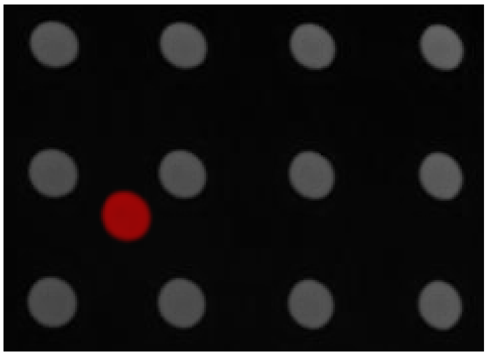

Figure 3: *Example of piecewise-linear point spread function interpolation. Gray kernels correspond to images of blurred point sources, whereas the red kernel is linearly interpolated from its three nearest neighbors.*

sources or printed test patterns (and corresponding assumptions on the form of the PSF), we observe that the out-of-focus image of a point light directly gives the point spread function. As a result, we propose using a two-dimensional array of point (white) light sources that can either be printed or projected from an absorbing (black) test surface.

3.1. Capture Setup

We display the test pattern shown in Figure 4(a) using a NEC MultiSync LCD (Model 2070NX). The calibration images were collected using a Canon EOS Rebel XT with a Canon 100mm Macro Lens. The lens was modified to allow the manual insertion of aperture patterns directly into the plane of the iris (i.e., by removing the original lens diaphragm). A typical calibration image, collected with an open aperture, is shown in Figure 4(b). (Note the characteristic cat's eye pattern.) To further illustrate the behavior of our modified lens, we have shown the calibration image acquired with a "star-shaped" aperture in Figure 4(c).

3.2. Parametric Model Estimation

Given the captured PSF data, we begin by segmenting the individual kernels using basic image processing and morphological operations. Next, we approximate the image-coordinate projection of a point light source as the intensity centroid of the corresponding PSF kernel. Finally, we approximate the local vignetting by averaging the values observed in each kernel. We proceed by interpolating the sparse set of vignetting coefficients using a low-order polynomial model. Similarly, we use a piecewise-linear interpolation scheme inspired by [NO98] to obtain a dense estimate of the spatially-varying PSF; first, we find the Delaunay triangulation of the PSF intensity centroids. For any given pixel, we linearly weight the PSF's on the vertices of the enclosing triangle using barycentric coordinates. Typical results are shown in Figure 3.

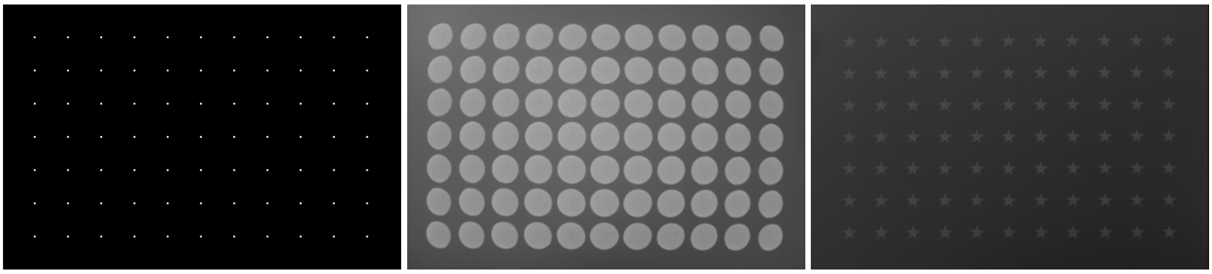

(a) calibration pattern (b) image with open aperture (small f-number) (c) image with "star-shaped" aperture

Figure 4: *Example of optical vignetting calibration using a two-dimensional array of point light sources. (a) The calibration image containing an array of* 7×11 *point light sources. (b) An image acquired with an open aperture that exhibits the characteristic "cat's eye" effect [Ray02]. (c) An image obtained by placing a mask with a "star-shaped" pattern in the aperture.*

4. Synthesis of Vignetting and Aperture Effects

Next, we focus on simulating the previously-discussed vignetting and aperture-dependent effects. In particular, we recall that controlling the bokeh is of particular importance to both professional and casual photographers. Through bokeh, the shape of out-of-focus points can be manipulated to impart additional meaning or stylization to an image. For example, as shown in Figure 5(a), a smooth bokeh can create an enhanced sense of separation between the foreground and background. This effect is typically exploited in portrait and macro photography, with certain lenses becoming prized in these fields for their exquisite bokeh. Similarly, distinct aperture shapes (e.g., hearts, stars, diamonds, etc.) can be used to for a novel effect or to convey a particular meaning (see Figure 5(b)). In the following sections we'll propose several methods for controlling the bokeh after image acquisition.

4.1. The *Bokeh Brush*

Recall that the bokeh is a direct result of the spatially-varying point spread function, which is itself due to the shape of the aperture (or other occluding structures in the lens). Traditionally, photographers would have to carefully select a lens or aperture filter to achieve the desired bokeh at the time of image acquisition. Inspired by computational photography, we present a novel solution for post-capture, spatially-varying bokeh adjustment.

(a) typical bokeh for portraits (b) "star-shaped" bokeh

Figure 5: *Illustration of bokeh in conventional photography.*

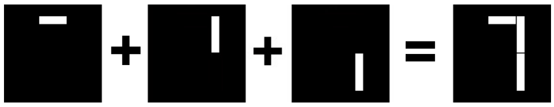

Figure 6: *Example of aperture superposition.*

4.1.1. Aperture Superposition Principle

Recall that, for unit magnification, the recorded image irradiance $I_i(x,y)$ at a pixel (x,y) is given by

$$I_i(x,y) = \iint_{\Omega} B(s,t;x,y) I_o(s,t)\, ds\, dt, \tag{1}$$

where Ω is the domain of the image, $I_o(x,y)$ is the irradiance distribution on the object plane, and $B(s,t;x,y)$ is the spatially-varying point spread function [HK07, NO98]. The PSF can also be expressed as a linear superposition of N basis functions $\{B_i(s,t;x,y)\}$ such that

$$I_i(x,y) = \sum_{i=1}^{N} \lambda_i \iint_{\Omega} B_i(s,t;x,y) I_o(s,t)\, ds\, dt. \tag{2}$$

This result indicates a direct and simple method to control bokeh in post-processing. Since the spatially-varying point spread function is dependent on the shape of the aperture, we find that rather than using only a single user-selected aperture, we can record a series of photographs using a small subset of *basis apertures* $\{A_i(s,t;x,y)\}$ that span a large family of iris patterns. As shown in Figure 6, a given aperture function $A(s,t;x,y)$ can then be approximated by the following linear combination.

$$A(s,t;x,y) = \sum_{i=1}^{N} \lambda_i A_i(s,t;x,y). \tag{3}$$

This expression states the *aperture superposition principle*: the images recorded with a given set of basis apertures can be linearly combined to synthesize the image that would be formed by the aperture resulting from the same combination.

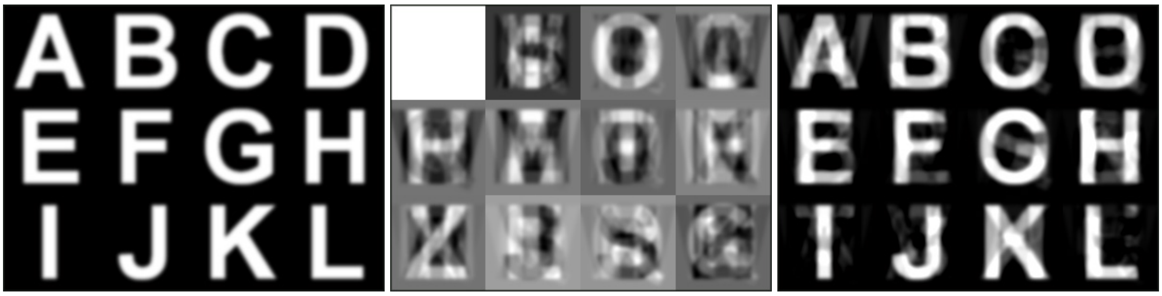

(a) training apertures (b) eigenaperture set (c) training aperture reconstructions

Figure 7: Bokeh Brush *PCA-derived apertures (spanning the capitalized* Arial *font). (a) A subset of 12 training apertures* $\{\mathbf{x}_i\}$*. (b) From left to right and top to bottom: The open aperture, the normalized offset aperture* $\tilde{\mathbf{x}}_0$*, and the first ten components of the PCA-derived eigenaperture set* $\{\tilde{\phi}_j\}$*. (c) The training aperture reconstructions* $\{\tilde{\mathbf{x}}_i\}$*.*

4.1.2. Bokeh Synthesis using Principal Components

Although we have demonstrated that images with different apertures can be linearly combined, we still require an efficient basis. One solution would be to use a set of translated pinholes; such a strategy would record the incident light field [LLC07]. While acknowledging the generality of this approach, we observe that specialized bases can be used to achieve greater compression ratios. In this section, we apply principal component analysis (PCA) to compress an application-specific set of apertures and achieve post-capture bokeh control without acquiring a complete light field.

Let's begin by reviewing the basic properties of PCA, as popularized by the *eigenfaces* method introduced by Turk and Pentland [TP91]. Assume that each d-pixel image is represented by a single $d \times 1$ column vector $\mathbf{x}_i$. Recall that the projection $\mathbf{x}'_i$ of $\mathbf{x}_i$ on a linear subspace is

$$\mathbf{x}'_i = \Phi^T(\mathbf{x}_i - \bar{\mathbf{x}}), \tag{4}$$

where Φ is a $d \times m$ matrix (with $m < d$), whose columns form an orthonormal basis for a linear subspace of $\mathbb{R}^d$ with dimension m. Also note that we have subtracted the mean image $\bar{\mathbf{x}} = \frac{1}{N}\sum_{i=1}^{N}\mathbf{x}_i$. For the particular case of PCA, the columns of Φ correspond to the first m unit-length eigenvectors $\{\phi_j\}$ (sorted by decreasing eigenvalue) of the $d \times d$ covariance matrix Σ given by

$$\Sigma = \frac{1}{N}\sum_{i=1}^{N}(\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})^T.$$

We refer to the m eigenvectors $\{\phi_j\}$ as the *principal components* of the data. The least-squares reconstruction $\hat{\mathbf{x}}_i$ of $\mathbf{x}_i$ is given by

$$\hat{\mathbf{x}}_i = \bar{\mathbf{x}} + \Phi\mathbf{x}'_i. \tag{5}$$

Now that we have reviewed the basic properties of PCA, let's use it to compress any given set of apertures. In post-processing, a photographer may want to select from a broad class of aperture shapes – ones which could vary from image to image or even within the same picture. For example, a novel application could include spanning the set of apertures corresponding to the capitalized letters in the Arial font (see Figure 7(a)). Note that the eigenvectors $\{\phi_j\}$ obtained by analyzing the set of non-negative "training" aperture images $\{\mathbf{x}_i\}$ will be signed functions on $\mathbb{R}^d$. Since we can only manufacture non-negative apertures for use with incoherent illumination, we will need to scale these eigenvectors. Let use define the set $\{\tilde{\phi}_j\}$ of d-dimensional real-valued *eigenapertures* on the range $[0,1]$ which satisfy

$$\tilde{\Phi} = (\Phi - \beta_1)\alpha_1{}^{-1},$$

where β_1 and α_1 are the necessary bias and scaling matrices, respectively. As before, we propose recording a sequence of images of a static scene using each individual eigenaperture. Afterwards, we can reconstruct the PCA-based estimate $\hat{\mathbf{I}}$ of an image $\mathbf{I}$ collected by any aperture function $\mathbf{x}$. We note that the best-possible aperture approximation $\hat{\mathbf{x}}$ is given by

$$\hat{\mathbf{x}} = \tilde{\Phi}\alpha_1\lambda + \beta_1\lambda + \mathbf{x}_0, \tag{6}$$

where the projection coefficients λ and the offset aperture $\mathbf{x}_0$ are given by

$$\lambda = \Phi^T\mathbf{x} \quad \text{and} \quad \mathbf{x}_0 = \bar{\mathbf{x}} - \Phi\Phi^T\bar{\mathbf{x}}. \tag{7}$$

Typical reconstruction results are shown in Figure 7(c). Since we cannot use a negative-valued offset mask, we further define the normalized offset aperture $\tilde{\mathbf{x}}_0$ such that

$$\tilde{\mathbf{x}}_0 = (\mathbf{x}_0 - \beta_2)\alpha_2{}^{-1}, \tag{8}$$

where β_2 and α_2 are the necessary bias and scaling terms, respectively. Combining Equations 6, 7, and 8 and assuming a spatially-invariant PSF, we conclude that the best reconstruction $\hat{\mathbf{I}}$ of an image $\mathbf{I}$ collected with the aperture function $\mathbf{x}$ is given by the following expression.

$$\hat{\mathbf{I}} = \mathbf{I} * \hat{\mathbf{x}} = \mathbf{I} * (\tilde{\Phi}\alpha_1\lambda) + \mathbf{I} * (\beta_1\lambda) + \alpha_2(\mathbf{I} * \tilde{\mathbf{x}}_0) + \beta_2\mathbf{I} \tag{9}$$

From this relation it is clear that $m+2$ exposures are required to reconstruct images using m eigenapertures – since images with open and normalized offset apertures must also be recorded. Note that a similar synthesis equation could also be used with spatially-varying point spread functions.

4.1.3. Bokeh Synthesis using Non-negative Factorization

As an alternative to eigenapertures, we propose applying non-negative matrix factorization (NMF) to the training apertures to directly obtain a non-negative basis [LS99]. As shown in Figure 8, the reconstruction from NMF-derived apertures is similar in quality to that obtained using PCA. Note that NMF eliminates the need for either an open aperture or bias aperture – reducing the number of required exposures for a given number of basis apertures when compared to PCA. Unfortunately, unlike PCA, the basis produced by our NMF implementation is not unique and will depend on the initial estimate of the non-negative factorization.

5. Results

5.1. Spatially-varying Deblurring

The procedure for estimating a spatially-varying PSF, as outlined in Section 3, was verified by simulation. As previously discussed, deconvolution using spatially-varying blur kernels has been a long-term topic of active research in the computer vision community [NO98, ÖTS94]. For this paper, we chose to implement a piecewise-linear PSF interpolation scheme inspired by the work of Nagy and O'Leary [NO98]. Typical deblurring results are shown in Figure 9.

5.2. Vignetting Synthesis

The *Bokeh Brush* was evaluated through physical experiments as well as simulations. As shown in Figure 10, a sample scene containing several point scatterers was recorded using a "seven-segment" aperture sequence; similar to the displays in many handheld calculators, the "seven-segment" sequence can be used to encode a coarse approximation of the Arabic numerals between zero and nine, yielding a compression ratio of 1.43. A synthetic "8" aperture was synthesized by adding together all the individual segment aperture images. Note that the resulting image is very similar to that obtained using an '8"-shaped aperture.

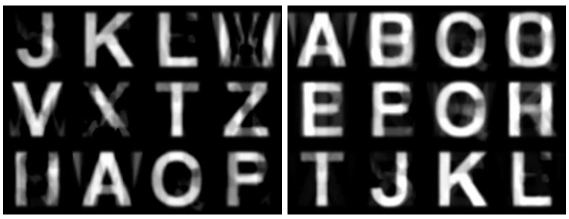

(a) NMF-derived apertures (b) approximation results

Figure 8: Bokeh Brush *NMF-derived apertures (spanning the capitalized* Arial *font). (a) First twelve basis apertures. (b) The resulting approximations of the training apertures.*

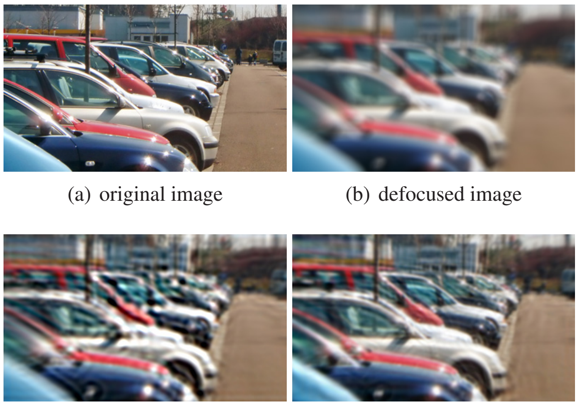

(a) original image (b) defocused image

(c) deblurred (mean PSF) (d) deblurred (spatially-varying)

Figure 9: *Example of deconvolution using a calibrated spatially-varying PSF. (a) The original image. (b) A simulated uniformly-defocused image. (c) Deconvolution results using the mean PSF. (d) Deconvolution results using the estimated spatially-varying PSF with the method of [NO98].*

The PCA-derived basis apertures initially proved difficult to manufacture – since they require precise high-quality printing processes. As an alternative, we confirm their basic design via simulation. As shown in Figure 11, a sample HDR scene was blurred using a spatially-varying PSF which is linearly proportional to depth. Note that this approximate depth-of-field effect has recently been applied to commercial image manipulation software, including Adobe's "lens blur" filter [RV07]. As shown in the included examples, the image synthesis formula given in Equation 9 was applied successfully to model novel aperture shapes. For this example, a total of 12 apertures were used to span the capitalized Arial characters, yielding a compression ratio of 2.17.

Finally, we note that the proposed method will also allow per-pixel bokeh adjustment. In particular, the individual reconstructions were interactively combined in Figure 11(c) in order to spell the word "BOKEH" along the left wall. We believe that such applications effectively demonstrate the unique capability of the *Bokeh Brush* to facilitate physically-accurate image stylization.

6. Discussion of Limitations

The primary limitation of our analysis and synthesis methods is that they neglect effects due to diffraction. In addition, the proposed *Bokeh Brush* will only work for static scenes, although one can imagine certain configurations with multiple cameras and beam-splitters to obtain real-time measurements. We recognize that using point light sources could be inefficient (versus line or area sources), since long exposures will be required. In addition, both the vignetting and PSF kernels are only available at discrete positions and must be interpolated to obtain per-pixel estimates. In the future, light

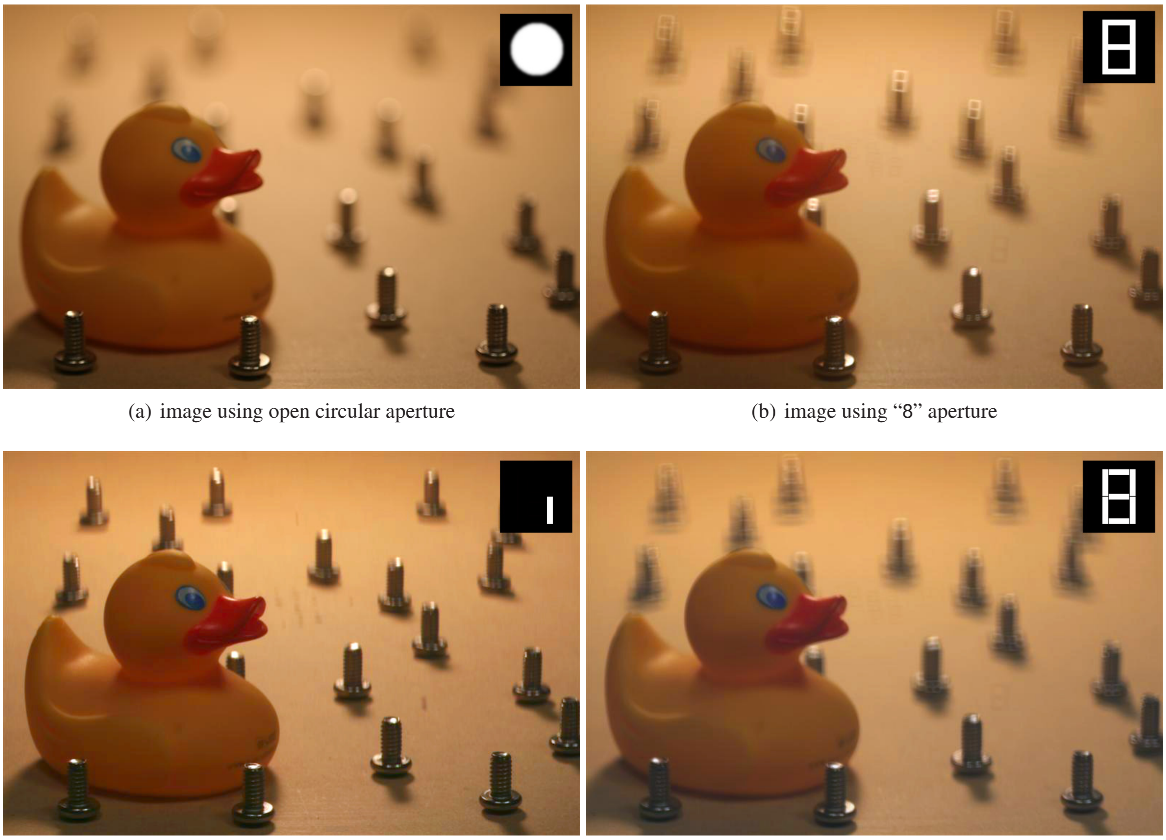

(a) image using open circular aperture

(b) image using “8” aperture

(c) image using second segment of “seven-segment” sequence

(d) reconstructed “8” aperture image

Figure 10: Bokeh Brush *experimental results for a “seven-segment” aperture sequence. (a) Image obtaining using an open aperture (with a small f-number). (b) Scene recorded by inserting an “8”-shaped aperture in the iris plane of a conventional lens. (c) Scene recorded by inserting a single segment in the iris plane. (d) Image reconstructed by aperture superposition (i.e., by summing the individual “seven-segment” aperture contributions via Equation 3).*

field cameras may become commonplace; in this situation, we recognize that compressed aperture bases would not be necessary.

7. Conclusion

We have analyzed optical vignetting in the context of methods in computational photography and have shown that it plays an important role in image formation. In particular, by exploiting the simple observation that the out-of-focus image of a point light directly gives the point spread function, we have shown a practical low-cost method to simultaneously estimate the vignetting and the spatially-varying point spread function. Similarly, we have shown the novel *Bokeh Brush* application which, to our knowledge, constitutes the first means of modifying the bokeh after image acquisition in an efficient and physically-accurate manner. Overall, we hope to inspire readers to think about vignetting and bokeh as expressive methods for enhancing the effects of depth-of-field, high intensity points, and aesthetics.

Acknowledgements

We would like to thank Martin Fuchs and Amit Agrawal for their helpful suggestions while the authors were at MERL. We would also like to thank the following Flickr members: Carlos Luis (for Figure 5(a)) and Harold Davis (for Figure 5(b) from http://www.digitalfieldguide.com/).

References

[AAB96] ASADA N., AMANO A., BABA M.: Photometric calibration of zoom lens systems. In *Proc. of the International Conference on Pattern Recognition* (1996), pp. 186–190.

[BD07] BAE S., DURAND F.: Defocus magnification. *Computer Graphics Forum 26*, 3 (2007).

[d'A07] D'ANGELO P.: Radiometric alignment and vignetting calibration. In *Camera Calibration Methods for Computer Vision Systems* (2007).

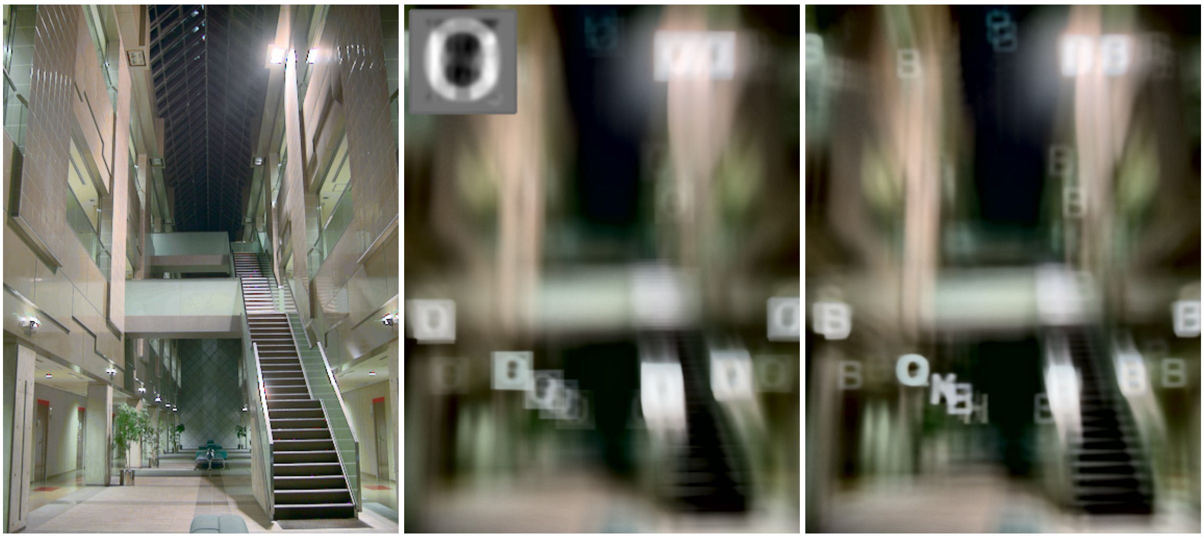

(a) original HDR image (b) simulated for first eigenaperture ($\tilde{\phi}_1$) (c) example of Bokeh stylization

Figure 11: Bokeh Brush *simulation results. (a) Input high dynamic range image. (b) Example of scene simulated using the first eigenaperture function and an approximated depth-of-field effect. (c) Example of Bokeh stylization where the aperture function has been adjusted in a spatially-varying manner to read "BOKEH" along the left wall.*

[GC05] GOLDMAN D. B., CHEN J.-H.: Vignette and exposure calibration and compensation. In *Proc. of the International Conference on Computer Vision* (2005), pp. 899–906.

[HK06] HASINOFF S. W., KUTULAKOS K. N.: Confocal stereo. In *Proc. of the European Conference on Computer Vision* (2006).

[HK07] HASINOFF S. W., KUTULAKOS K. N.: A layer-based restoration framework for variable-aperture photography. In *Proc. of the International Conference on Computer Vision* (2007).

[KW00] KANG S. B., WEISS R. S.: Can we calibrate a camera using an image of a flat, textureless lambertian surface? In *Proc. of the European Conference on Computer Vision* (2000), pp. 640–653.

[LFDF07] LEVIN A., FERGUS R., DURAND F., FREEMAN W. T.: Image and depth from a conventional camera with a coded aperture. *ACM Trans. Graph. 26*, 3 (2007).

[LLC07] LIANG C.-K., LIU G., CHEN H.: Light field acquisition using programmable aperture camera. In *Proc. of the International Conference on Image Processing* (2007).

[LS99] LEE D. D., SEUNG H. S.: Learning the parts of objects by non-negative matrix factorization. *Nature 401*, 6755 (October 1999), 788–791.

[LS05] LITVINOV A., SCHECHNER Y. Y.: Addressing radiometric nonidealities: A unified framework. *Proc. of International Conference on Computer Vision and Pattern Recognition* (2005), 52–59.

[NO98] NAGY J. G., O'LEARY D. P.: Restoring images degraded by spatially variant blur. *SIAM J. Sci. Comput. 19*, 4 (1998), 1063–1082.

[ÖTS94] ÖZKAN M. K., TEKALP A. M., SEZAN M. I.: POCS-based restoration of space-varying blurred images. *IEEE Transactions on Image Processing 3*, 4 (1994).

[Ray02] RAY S. F.: *Applied Photographic Optics*. Focal Press, 2002.

[RV07] ROSENMAN R., VICANEK M.: Depth of Field Generator PRO, 2007. http://www.dofpro.com.

[TP91] TURK M. A., PENTLAND A. P.: Face recognition using eigenfaces. In *Proc. of the International Conference on Computer Vision and Pattern Recognition* (1991).

[VRA*07] VEERARAGHAVAN A., RASKAR R., AGRAWAL A., MOHAN A., TUMBLIN J.: Dappled photography: Mask enhanced cameras for heterodyned light fields and coded aperture refocusing. *ACM Trans. Graph. 26*, 3 (2007), 69.

[vW07] VAN WALREE P.: Vignetting, 2007. http://www.vanwalree.com/optics/vignetting.html.

[Yu04] YU W.: Practical anti-vignetting methods for digital cameras. *IEEE Transactions on Consumer Electronics* (2004), 975–983.

[ZLK06] ZHENG Y., LIN S., KANG S. B.: Single-image vignetting correction. In *Proc. of International Conference on Computer Vision and Pattern Recognition* (2006), pp. 461–468.