

Utilizing Optical Aberrations for Extended-Depth-of-Field Panoramas

Huixuan Tang and Kiriakos N. Kutulakos*

Dept. of Computer Science, University of Toronto
{hxtang, kyros}@cs.toronto.edu

Abstract. Optical aberrations in off-the-shelf photographic lenses are commonly treated as unwanted artifacts that degrade image quality. In this paper we argue that such aberrations can be useful, as they often produce point-spread functions (PSFs) that have greater frequency-preserving abilities in the presence of defocus compared to those of an ideal thin lens. Specifically, aberrated and defocused PSFs often contain sharp, edge-like structures that vary with depth and image position, and that become increasingly anisotropic away from the image center. In such cases, defocus blur varies spatially and preserves high spatial frequencies in some directions but not others. Here we take advantage of this fact to create extended-depth-of-field panoramas from a set of overlapping photos taken with off-the-shelf lenses and a wide aperture. We achieve this by first measuring the lens PSF through a one-time calibration procedure and then using multi-image deconvolution to restore anisotropic blur in areas of image overlap. Our results suggest that common wide-aperture lenses may preserve frequencies well enough to allow extended-depth-of-field panoramic photography with large apertures, resulting in potentially much shorter exposures.

1 Introduction

Optical aberrations—deviations of a lens system from the predictions of paraxial optics—occur in all photographic lenses as the inevitable compromise between image quality, lens complexity and cost. Aberrations affect image quality to various degrees and become especially significant when capturing photos with a wide lens aperture [1].

Optical aberrations are typically considered as undesirable artifacts that cause well-focused subjects to appear blurry in a photo. Although recent work has shown that it is sometimes possible to restore the appearance of in-focus subjects by undoing this blur [2], the impact of lens aberrations on a photo’s *out-of-focus* regions has not been explored by the vision community. Here we argue that aberrations and defocus in off-the-shelf lenses interact in ways that can have a significant impact on their point spread function (PSF): as shown in Figs. 1 and 2b, aberrations can introduce sharply-defined edges in the PSF of defocused points, making them non-uniform, anisotropic, and spatially varying. This represents a significant departure from the spatially-invariant PSF of an ideal thin lens, which is always a uniform-intensity disk (*i.e.*, a “pillbox function”).

* This work was supported by the Natural Sciences and Engineering Research Council of Canada under the Discovery, Discovery Accelerator, GRAND-NCE and RGPIN programs.




Fig. 1: The PSF of a Canon 50mm $f/1.2L$ lens focused at $2m$ for a plane at depth $1.5m$. Note the PSF’s spatially-varying and strongly-anisotropic structure.

A key consequence of this structure is that aberrated PSFs have a greater ability to preserve high spatial frequencies in regions of a scene that are out of focus. We empirically analyze the frequency-preserving behavior of such PSFs using a simplified lens aberration model that takes into account Seidel aberrations up to third order. More specifically, we pay attention to three properties of aberrated and defocused PSFs: (1) the way they vary in radial directions away from the image center, (2) their anisotropic structure near the image border and (3) asymmetries in how they vary as functions of object depth relative to the in-focus plane.

We argue that these properties of aberrated PSFs are useful for light-efficient acquisition of extended-depth-of-field panoramas. In particular, instead of capturing and stitching photos with a small aperture—which yields photos with a wide depth of field but requires long exposure times because of light inefficiency—we capture wide-aperture photos using short exposures. We then extend the panorama’s depth of field by restoring out-of-focus blur. This is possible because aberrations cause out-of-focus scene points to be blurred differently depending on their position on the image plane. Hence, in areas of image overlap, where the same scene point is blurred differently in each photo, aberrations can preserve frequencies well enough in all directions to enable significant deblurring and depth-of-field extension.

Our approach is related to four lines of recent work. First, work on aberration modeling [1, 3], PSF estimation [4], and aberration correction of in-focus subjects [2, 5] has also noted the spatially-varying nature of real-lens PSFs. Unlike all this work, our emphasis here is on studying the structure of aberrated *and defocused* PSFs for the purpose of counteracting defocus. Second, our work is similar in spirit to techniques that use specialized optics for depth-of-field extension [6, 7] but the PSFs we study apply to standard photographic lenses and thus have different characteristics. For instance, we observe that aberrated PSFs preserve high spatial frequencies only for object depths on one side of the in-focus plane (*i.e.*, either closer to or farther from the in-focus plane, but not both). Third, deblurring from multiple images with complementary PSFs has been studied in the context of coded-aperture photography and motion deblurring [8] and [9]. Here we use the same underlying principle but in a completely different imaging do-

main, where complementary PSFs come from aberrations rather than coded apertures or camera motions. Last but not least, our work can be thought of as a generalized mosaicing procedure [10], where scene points are imaged under different conditions in a dense collection of photos captured while panning a specially-equipped camera. Instead of relying on specialized hardware to independently modulate the appearance of each scene point, here we exploit aberrations, which are inherent in ordinary camera lenses. Moreover, our approach requires a relatively small number of images: just like any standard panorama construction method [11], there must be sufficient overlap between photos to enable registration and to ensure that individual scene points are imaged in at least two of them.

2 Modeling Aberrated Lens PSFs

In this section we consider the problem of modeling and estimating the non-stationary, depth-varying PSF of real camera lenses. We first measure in a controlled lab setting the “ground truth” PSF of a lens for different defocus levels and then propose a parametric model that behaves consistently with our acquired data.

2.1 PSF Measurement

Our measurement procedure seeks to acquire a dense 3D grid of “local” 2D PSFs that capture PSF variation over the image plane as well as over object depth (*i.e.*, defocus level). A 2D grid of local PSFs corresponding to a single depth is shown in Fig. 1.


To estimate spatially-varying PSFs at a specific object depth, we use the deconvolution-based approach of Joshi *et al.* [4]. The idea is to capture an image of a fronto-parallel plane containing a known pattern, use the pattern’s corners to align it to the image, and then estimate local PSFs using non-blind deconvolution. We decided to use this approach because it allows us to compute a dense grid of local PSFs from a single image and because it delivers PSFs of reasonable quality.¹

Since precise localization of corners is difficult for large levels of defocus (*e.g.*, PSF width > 50 pixels), we capture a second photo at each object depth with a narrow $f/16$ aperture and use this photo for image-to-pattern alignment (Fig. 2a). To capture the full 3D grid of local PSFs, we fix the zoom and focus settings of the camera and translate it along its optical axis using a translation stage to vary the object-to-camera distance.

We applied this procedure to three Canon lenses: a 50mm $f/1.2L$, an 85mm $f/1.2L$ and a 24-70mm $f/2.8$ lens with its zoom setting set to 70mm. In all cases, we used the widest-possible lens aperture and acquired a 19×13 grid of local PSFs for each of 20 to 60 object depths away from the in-focus plane, until the PSF’s diameter was approximately 80 pixels.² The defocused PSFs of all three lenses have evident non-uniform, edge-like features, with significant differences between center and corner PSFs (Fig. 2b). The depth ranges where these features appear, however, are specific to each

¹ We also tried to directly acquire the lens PSF by taking photos of pinhole arrays but this approach is prone to noise for large levels of defocus and can be affected by diffraction.

² This is approximately $0.5mm$ on the sensor of the Canon 1Ds Mk3 we used for acquisition.



(a) Patches from a narrow/wide-aperture image pair. (b) Measured PSFs for three Canon lenses at two positions on the image plane.

Fig. 2: Measuring the lens PSF for object distances outside the depth of field. For the PSFs in (b), lenses were focused at infinity, $0.95m$ and $0.38m$, respectively, with objects at a distance of $2m$, $1.2m$ and $0.5m$.

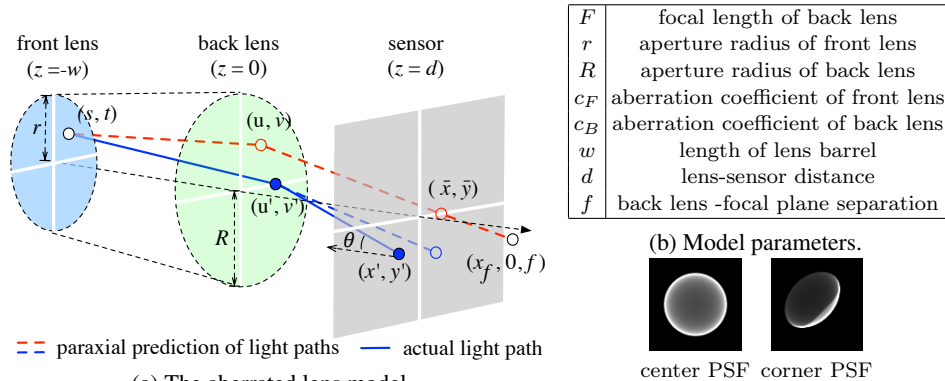


Fig. 3: Accounting for lens aberrations with a variable-cone doublet model.

lens: the 50mm lens yields defocused PSFs with sharp features when the lens is focused at infinity whereas the other two lenses yield such PSFs for object depths that are more distant than the in-focus plane. It follows that to take full advantage of the lenses’ frequency-preserving properties, it is important to choose a focus setting that places all objects of interest on the “frequency-preserving side” of the in-focus plane.

2.2 Modeling Lens Aberrations

The structure exhibited by our measured PSFs is primarily due to monochromatic aberrations and vignetting. Monochromatic aberrations are deviations from the paraxial lens approximation that warp the light paths passing through the lens.³ Vignetting causes a

³ Real lenses also suffer from chromatic aberrations [12]. These aberrations occur because focus and magnification depend on the wavelength of incident light. Although we do not consider such aberrations here, they can be taken into account by capturing and modeling local lens PSFs separately for each color channel.

reduction in brightness at the periphery of the image and is due to physical ray occlusion from the lens mechanics as well as natural light fall-off from oblique rays.

We now consider a simplified lens model that captures both monochromatic aberrations and vignetting. Our model consists of two thin lenses enveloped in a variable-cone-shaped barrel (Fig. 3a) and its parameters are summarized in Fig. 3b. To model aberrations, we assume that each lens is rotationally symmetric so that deviations from the paraxial approximation of a ray can be expressed relative to the *meridional plane*, which contains both the paraxial ray and the optical axis. We also restrict our model to third-order optics, where the angle of a refracted ray is expressed as a cubic polynomial of the ray's height at the lens interface. To model vignetting, we use the variable-cone aperture model of Asada *et al.* [13] in which ray occlusion (and thus vignetting) is not affected by aberrations.

Fig. 3c shows example defocused PSFs produced by our lens model. These PSFs are qualitatively very similar to the captured PSFs shown in Fig. 2b. This suggests that our simplified model is rich enough to capture the overall structure of aberrated PSFs that we observe in practice.

We now sketch the derivation of our PSF model; a detailed derivation can be found in the supplementary materials [14]. By definition, a local PSF is the image of an idealized isotropic point light source. This image is determined by the cone of rays exiting the source. We parameterize these rays by their intersections (u, v) with the back lens (Fig. 3a) and assume, without loss of generality, that the point source forms a focused image at $(x_f, 0, f)$ under the paraxial approximation. Together, (u, v) and $(x_f, 0, f)$ completely determine the ray's path, $(s, t) \rightarrow (u, v) \rightarrow (\bar{x}, \bar{y})$, under paraxial optics as well as its path $(s, t) \rightarrow (u', v') \rightarrow (x', y')$ under our third-order model.

It follows that the local 2D PSF $k(x, y)$ can be expressed as an irradiance integral of aberrated rays

$$k(x, y) = \int_{u^2+v^2 \leq R^2} \delta(x - x'(u, v), y - y'(u, v)) p(u, v) dudv \quad (1)$$

where $\delta(\cdot)$ denotes Dirac's delta and $p(u, v)$ is the *pupil function* which models vignetting. This function is zero if the corresponding ray is occluded by the lens and is $\cos^4 \theta$ otherwise, where θ is the angle between the ray and the normal of the sensor plane, *i.e.*, $\theta = \cos^{-1} \left(f / \sqrt{(x_f - u)^2 + v^2 + f^2} \right)$.

From the thin-lens law, the intersection (s, t) of the paraxial ray and the front lens is

$$\begin{pmatrix} s \\ t \end{pmatrix} = m_P \begin{pmatrix} u \\ v \end{pmatrix} + \begin{pmatrix} s_0 \\ 0 \end{pmatrix}, \text{ where } m_P = w \left(\frac{1}{w} + \frac{1}{f} - \frac{1}{F} \right) \text{ and } s_0 = \frac{-x_f w}{f}. \quad (2)$$

Restricting (s, t) to its circular aperture determines the pupil function in Eq. (1):

$$p(u, v) = \begin{cases} \cos^4 \theta, & \text{if } (m_P u + s_0)^2 + (m_P v)^2 \leq r^2 \text{ and } u^2 + v^2 \leq R^2 \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

We can now derive the ray displacement on the image plane using Eq. (2):

$$\begin{aligned} \begin{pmatrix} \Delta x(u, v) \\ \Delta y(u, v) \end{pmatrix} &\doteq \begin{pmatrix} x' \\ y' \end{pmatrix} - \begin{pmatrix} \bar{x} \\ \bar{y} \end{pmatrix} = \underbrace{\alpha_1(u^2 + v^2)}_{\text{spherical}} \begin{pmatrix} u \\ v \end{pmatrix} + \underbrace{\alpha_2 x_f \begin{pmatrix} 3u^2 + v^2 \\ 2uv \end{pmatrix}}_{\text{coma}} \\ &+ \underbrace{\alpha_3 x_f^2 \begin{pmatrix} u \\ 0 \end{pmatrix}}_{\text{astigmatism}} + \underbrace{\alpha_4 x_f^2 \begin{pmatrix} u \\ v \end{pmatrix}}_{\text{field curvature}} + \underbrace{\alpha_5 x_f^3 \begin{pmatrix} 1 \\ 0 \end{pmatrix}}_{\text{field distortion}} \end{aligned} \quad (4)$$

where

$$\begin{aligned} \alpha_1 &= d \left(c_F \frac{F-w}{F} m_P^3 + c_B \right), \quad \alpha_2 = -c_F \left(w + d \frac{F-w}{F} \right) m_P^2 \left(\frac{w}{f} \right), \\ \alpha_3 &= 2c_F \left(w + d \frac{F-w}{F} \right) m_P \left(\frac{w}{f} \right)^2, \quad \alpha_4 = c_F \left(w + d \frac{F-w}{F} \right) m_P \left(\frac{w}{f} \right)^2 \\ \text{and } \alpha_5 &= -c_F \left(w + d \frac{F-w}{F} \right) \left(\frac{w}{f} \right)^3. \end{aligned} \quad (5)$$

Eq. (4) provides a compact model of ray displacements as a linear combination of five displacement vector fields. These fields are referred to as the primary Seidel aberrations [1]. The five fields approximate the law of refraction up to third order and are usually good enough to model both spherical and aspherical lenses used in contemporary commercial lenses. When the lens settings are fixed and objects lie on a single depth plane, it suffices to use Eq. (4) to describe monochromatic aberrations by treating the parameters $\alpha_1, \dots, \alpha_5$ as constants. However, Eq. (5) is useful if multiple focus-settings and object depths are involved because it explicitly models the displacement fields' dependence on depth. Note that both m_P and s_0 depend on $1/f$, the inverse focal distance, which depends on object depth. This factor can significantly affect the displacement of aberrated rays and explains why the PSF's structure changes drastically with depth.

3 Frequency-Preserving Properties of Aberrated PSFs

We now use the lens model of Section 2.2 to examine the frequency-preserving properties of aberrated PSFs. To do this, we employ a light field analysis similar to that of Levin *et al.* [15] and Zhang and Levoy [16].

Let us consider again the image of a point source whose paraxial image is at $(x_f, 0, f)$. We parameterize rays incident on the sensor by their intersection with two reference planes, one aligned with the back lens and the other aligned with the in-focus plane. For simplicity, we take the origin of the in-focus plane to be at $(x_f, 0, f)$. Paraxial optics predict that all rays from the source converge at $(x_f, 0, f)$. It follows that the light field incident on the sensor under paraxial optics is given by

$$l_0(x, y, u, v) = \delta(x, y) p(u, v). \quad (6)$$

Note that the coordinates (x, y) correspond to points on the in-focus plane and should not be confused with the image-plane coordinates in Eq. (1).

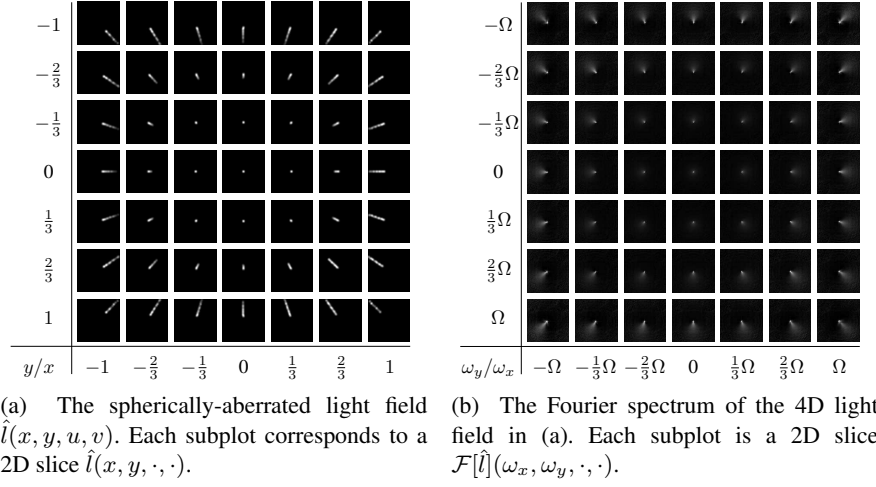


Fig. 4: The 4D light field due to an isotropic point source under our aberrated-lens model.

In the presence of aberrations, the light field is distorted according to

$$l(x, y, u, v) = \delta\left(x - \Delta x_f(u, v), y - \Delta y_f(u, v)\right) p(u, v) \quad (7)$$

where $(\Delta x_f, \Delta y_f)$ are obtained from Eqs. (4) and (5) by setting $d = f$.

From the Generalized Fourier Slice Theorem [17] it follows that the Fourier transform of the lens PSF, $k(\cdot)$, is a slice of the 4D Fourier transform of the light field:

$$\mathcal{F}[k](\mu, \nu) = \mathcal{F}[l](\alpha_f \mu, \alpha_f \nu, (1 - \alpha_f)\mu, (1 - \alpha_f)\nu), \quad \text{where } \alpha_f = d/f. \quad (8)$$

Now, let $\hat{l}(x, y, u, v) = \delta\left(x - \Delta x_f(u, v), y - \Delta y_f(u, v)\right)$ be the light field obtained by ignoring the effects of vignetting. From the convolution theorem, we have

$$\mathcal{F}[l](\omega_x, \omega_y, \omega_u, \omega_v) = \mathcal{F}[\hat{l}](\omega_x, \omega_y, \omega_u, \omega_v) \otimes \mathcal{F}[p](\omega_u, \omega_v) \quad (9)$$

where \otimes denotes convolution.

Eq. (9) tells us that the Fourier transform of the lens PSF is a convolution of two light fields—one that depends only on aberrations and one that depends only on vignetting. In the following, we first discuss how the five Seidel aberrations affect the aberrated light field and then discuss the effects of vignetting.

The *spherical aberration term*, $\alpha_1(u^2 + v^2)\left(\frac{u}{v}\right)$, causes rays passing through the boundary of the lens pupil to converge at a different depth compared to rays passing through a circle near the pupil's center. These displacements, which are independent of the local PSF's position on the image plane, cause sharp features in the PSF and break the PSF's symmetry relative to the in-focus plane. An example of a spherically-aberrated light field and its Fourier transform is shown in Fig. 4. From Eq. (8), it follows

that the region in 4D frequency space that contributes to the PSF is the domain of $\mathcal{F}[\hat{l}](\alpha_f \mu, \alpha_f \nu, (1 - \alpha_f) \mu, (1 - \alpha_f) \nu)$. Since the aberrated light field is non-zero outside of this domain, it wastes some energy outside of the “focal manifold” and therefore does not have the extended-depth-of-field capabilities of some recent computational cameras.⁴ Nevertheless, spherically-aberrated light fields do concentrate energy near the focal manifold and can preserve high frequencies even when the scene is not on the in-focus plane. Importantly, observe that the Fourier spectrum in Fig. 4b is large only for frequencies ω_x, ω_u (and ω_y, ω_v) having the same sign, *i.e.*, when $d < d_f$. This indicates that high frequencies are only preserved at object depths one side of the in-focus plane.

The coma aberration term, $\alpha_2 x_f \left(\frac{3u^2 + v^2}{2uv} \right)$, can be thought of as a change in magnification that depends on the ray’s position on the pupil plane, *i.e.*, the u - v plane. Thus, rays passing through the periphery of the pupil may cross the image plane at different distances from the paraxial image of the point source. This causes asymmetries in the local PSF that vary radially away from the image center. Coma aberrations produce a cubic phase delay in the optical wavefront, resulting in defocused PSFs that preserve high spatial frequencies in a way analogous to wavefront coding [6, 7].

Both *the astigmatism term, $\alpha_3 x_f^2 \left(\frac{u}{v} \right)$,* and *the field curvature term, $\alpha_4 x_f^2 \left(\frac{u}{v} \right)$,* correspond to displacements that are linear functions of the ray’s position on the pupil. Since ray displacements due to defocus are also linear functions of pupil position, these aberrations can be differentiated from defocus only by the fact that they increase quadratically with distance from the image center. In addition, astigmatism displaces rays only in the radial direction, yielding elliptically-shaped PSFs near the image corners.

The field distortion term, $\alpha_5 x_f^3 \left(\frac{1}{v} \right)$, causes a non-linear, radially-symmetric image distortion. The term, however, does not affect the local 2D PSF because it is completely independent of the pupil position. It can therefore be safely ignored when analyzing interactions between aberrations and defocus.

Effect of the pupil function. Eq. (3) suggests that the pupil function has a “cat eye”-shaped support on the pupil plane, with its shortest axis aligned with the v -axis (Fig. 5a). This causes its Fourier spectrum to be elongated along that axis. Since the pupil function $p(u, v)$ is a 2D function that is independent of (x, y) , its Fourier transform in the 4D light field space will have its spectrum concentrated on the $\mathcal{F}[p](0, 0, \omega_u, \omega_v)$ slice. Therefore, frequency-domain blurring due to the pupil function occurs only within each 2D subplot in Fig. 4b. An example of such a frequency-domain blur is shown Figs. 5a and 5b.

In summary, we observe that (1) when defocus is present, spherical and coma aberrations yield frequency-preserving PSFs; (2) this occurs only for depths on one side of the in-focus plane and thus depth-of-field extension—although possible—is asymmetric relative to that plane; (3) astigmatism and field curvature modulate defocus in a spatially-varying manner but do not affect the PSF’s frequency-preserving properties; and (4) vignetting makes the PSF even more anisotropic, with frequencies in the radial direction (*i.e.*, along the v -axis) preserved more than others.

⁴ See [15] for a thorough discussion.

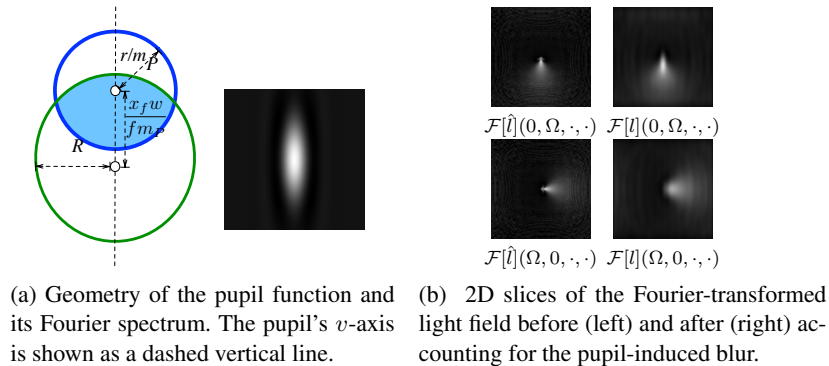


Fig. 5: The effect of vignetting. The Fourier-transformed pupil function in (a) is convolved with the individual 2D slices shown in Fig. 4b. Two of those slices are shown in the left column of (b). Note that vignetting blurs these slices primarily in the vertical direction.

4 Building Light-Efficient Panoramas

We now take advantage of these properties to create extended-depth-of-field panoramas in a “light-efficient” way, *i.e.*, using photos taken with large apertures and short exposure times. In particular, wherever these photos overlap, the underlying sharp image will be blurred by several defocused PSFs, each of which preserves high spatial frequencies in some directions but not others. By combining these photos we can therefore restore spatial frequencies in many directions, extending the depth of field of the final panorama.

Our restoration method includes several steps. First, we calibrate the lens by recovering a 3D grid of local PSFs, as explained in Section 2.1. We then capture a sequence of photos using a wide aperture and a focus setting that places the scene of interest on the frequency-preserving side of the in-focus plane. These photos are aligned geometrically by estimating pairwise homographies with Autostitch [11]. Finally, we use the multi-image restoration procedure described below to compute the panorama.

Because the PSF is spatially varying, we restore the panorama patch by patch, assuming that (1) each patch may contain objects at multiple depths and (2) for a given depth, the PSF is spatially-invariant within each patch. We restore individual patches using joint, non-blind deblurring using the pre-calibrated local 2D lens PSFs. Since these PSFs vary with depth, we run the restoration procedure once for each depth hypothesis and then use the restoration results across all hypotheses to compute a per-pixel depth map for each patch.

Let ψ be a patch in the underlying sharp panorama and let $\varphi_1, \varphi_2, \dots, \varphi_N$ be the patches corresponding to ψ in N overlapping photos. We assume that each of these patches is formed by blurring the hidden patch ψ with a depth-dependent PSF that is specific to each photo and has been pre-warped by the homography that maps photo pixels to panorama pixels. When the entire patch contains scene points at just one depth,

the observed patch is given by

$$\varphi_j = k_{\lambda^*}^j \otimes \psi + n, \quad (10)$$

where $k_{\lambda^*}^j$ is the pre-warped PSF corresponding to the j -th photo, λ^* is the true depth of the patch and n denotes Gaussian noise of variance η^2 . When the patch contains points at multiple depths, Eq. (10) generalizes to a layered model, where each layer’s appearance is described by this equation.

Under a Gaussian image prior and a depth hypothesis λ , we can obtain an estimate $\tilde{\psi}_\lambda$ of the hidden patch using joint Wiener deconvolution:

$$\tilde{\psi}_\lambda = \mathcal{F}^{-1} \left[\frac{1}{\eta^2} \sum_j \left(\mathcal{F}[\varphi_j] \overline{\mathcal{F}[k_\lambda^j]} \right) V_\lambda^{-1} \right] \text{ where } V_\lambda^{-1} = \frac{1}{\eta^2} \sum_j \|\mathcal{F}[k_\lambda^j]\|^2 + S^{-1}. \quad (11)$$

Here S is the variance of $\mathcal{F}[\psi]$ and $\overline{\mathcal{F}[\cdot]}$ denotes the complex conjugate of $\mathcal{F}[\cdot]$.⁵

To assign a depth λ to each pixel p within a patch, we construct a piece-wise smooth depth map using a Markov random field approach [18]. We use per-pixel reconstruction error as the data term

$$E_\lambda(p) = \sum_j \left(\varphi_j(p) - [k_d^j \otimes \tilde{\psi}_\lambda](p) \right)^2 \quad (12)$$

and use the $L1$ -norm between neighboring depths as the smoothness term (with a weight of 0.1). To synthesize the final panorama $\tilde{\psi}$ from the computed depth map $\lambda(p)$, we simply copy pixels from the restored patch at the optimal depth: $\tilde{\psi}(p) = \tilde{\psi}_{\lambda(p)}(p)$.

5 Experiments


5.1 PSF Evaluation

We start by considering the advantage conferred by aberrated PSFs over the standard pillbox PSF. In order to evaluate the extended-depth-of-field performance regardless of the scene, we use two criteria: (1) the expected restoration error over the distribution of hidden images and (2) the power of PSFs to discriminate between different depths.

According to Hasinoff *et al.* [19], the expected mean-squared error of the restored patch is given by

$$\mathbb{E} \left[\|\psi(p) - \tilde{\psi}(p)\|^2 \right] = \sum_{\mu, \nu} V_\lambda(\mu, \nu), \quad (13)$$

where $V_\lambda(\mu, \nu)$ is defined in Eq. (11). Fig. 6a plots this term as a function of PSF size for a local 2D PSF estimated from our Canon 50mm f/1.2L lens, and for the pillbox PSF at the same defocus level. It shows that the expected reconstruction error of the pillbox PSF grows drastically with defocus, whereas the error curve for the real-lens PSF remains flat near the in-focus position. In addition, reconstruction error increases at a slower rate for object depths behind the in-focus plane. This is consistent with our



(a) Expected mean squared restoration error as a function of PSF width. The curves represent an average over nine randomly-chosen image locations. (b) KL divergence between images of different defocus level, computed from PSFs at nine randomly-chosen locations.

Fig. 6: Comparative evaluation of aberrated and pillbox PSFs.

observation in Section 3, that the frequency-preserving ability of aberrated lenses is generally asymmetric relative to the in-focus plane.

To evaluate the PSF’s depth-discrimination power, we calculate the Kullback-Leibler (KL) divergence between images at different defocus levels, for both the aberrated and the pillbox PSF. A high KL divergence is desirable because it indicates a smaller chance of incorrect depth estimation. As we show in the supplemental materials [14], the KL divergence between images defocused by k_1^j and k_2^j corresponding to object depths λ_1 and λ_2 , respectively, is given by

$$\text{KL}(\lambda_1, \lambda_2) = \frac{1}{2} \sum_{\mu, \nu} \left[\log \left(\frac{\eta^2 + S \sum_j |\mathcal{F}[k_2^j]|^2}{\eta^2 + S \sum_j |\mathcal{F}[k_1^j]|^2} \right) + \frac{S}{\eta^2} \sum_j |\mathcal{F}[k_2^j]|^2 - \frac{\sum_j |\mathcal{F}[k_1^j]|^2}{\sum_j (\eta^2/S + |\mathcal{F}[k_1^j]|^2)} - \frac{S}{\eta^2} \frac{|\sum_j \mathcal{F}[k_2^j] \mathcal{F}[k_1^j]|^2}{\sum_j (\eta^2/S + |\mathcal{F}[k_1^j]|^2)} \right] (\mu, \nu) \quad (14)$$

Fig. 6b plots the KL divergence for our aberrated lens and for an aberration-free lens. The figure shows that the KL divergence is higher in the aberrated case, suggesting that depth recovery is easier in presence of aberrations. This is because the PSF’s frequency spectrum varies inhomogeneously with depth, resulting in a smaller correlation between blurred object textures at different depths. The figure also shows that unlike the pillbox PSF, where it is impossible to tell if a defocused scene point is in front or behind the in-focus plane, aberrated PSFs do not suffer from this depth-reflection ambiguity.

5.2 Panorama Restoration

Simulation We synthesized ten images of a scene containing three depth layers (Fig. 7). All images are focused at the same depth, and each image is associated with a specific, randomly-chosen PSF from our Canon 50mm 1.2L lens calibration data. The PSF sizes ranged from 20 to 60 pixels. We restored the scene with the algorithm discussed in Sec. 4. Despite the challenging defocus blur, our algorithm successfully restored the details in the image and recovered an approximate depth map.

⁵ We set $S = 1$ and $\eta^{-2} = 1e - 4$ in our implementation.




Fig. 7: Multi-image restoration on simulated images.

Real data We captured three datasets under different imaging conditions, from macro to landscape. We used the Canon 50mm $f/1.2L$ lens to capture the indoor “magazine” scene (Fig. 8) and the outdoor “portrait” scene (Fig. 9). For both datasets we focused the lens at the far end of the scene’s depth range so that the aberrated PSF preserves frequencies. We used the Canon 24-70mm $f/2.8L$ zoom lens to capture the close-up “static” scene (Fig. 10). In this case, the lens was focused at $0.38m$ with its focal length set to $70mm$. We always captured photos at the largest possible aperture ($f/1.2$ or $f/2.8$).

Individual photos are significantly contaminated by defocus in all three examples. Despite this, multi-image restoration was able to successfully recover scene details outside the original depth of field (*e.g.*, text and facial features) and to obtain a reasonable depth segmentation. For the magazine scene, we also show patch restoration results from a single photo. Although some details are recovered even in this case, multi-image restoration is of much higher quality.

In both the portrait and static scenes, restoration suffers from ringing artifacts. This is because our restoration procedure is rather simple and is not designed to handle depth discontinuities. In addition, scenes with significant depth variations cannot be aligned precisely by estimating a single homography for each photo. These issues may be resolved by incorporating multi-view geometry into the image formation model and by relying on more advanced deconvolution algorithms. Nevertheless, our restoration method recovers many details not visible in the directly-stitched panorama.

6 Conclusion

Our goal in this paper is to show that aberrations in many photographic lenses can be *significant* and *useful* in analyzing defocus under certain conditions. More specifically, we believe that optical aberrations possess greater ability to preserve high frequencies and can be used for depth-of-field extension as well as depth estimation. Moreover, these frequency-preserving properties exhibit several anisotropies, both on the image plane and across depth, that must be taken into account for accurate results.

We are currently in the process of analyzing the aberration properties of several off-the-shelf photographic lenses and are exploring several directions for future work. These include (1) studying the general depth-from-defocus problem for real, aberrated lenses, (2) estimating the aberration properties of a lens from a single photo without prior calibration, and (3) developing a unified framework for one-shot depth-from-defocus and blind image deblurring.

References

1. Smith, W.J.: *Modern Optical Engineering: The Design of Optical Systems*. McGraw-Hill Professional (2000)
2. Schuler, C., Hirsch, M., Harmeling, S., Scholkopf, B.: Non-stationary correction of optical aberrations. In: *Proc. IEEE Int. Conf. on Computer Vision*. (2011) 659–666
3. Janssen, A.J.E.M.: Extended nijboer-zernike approach for the computation of optical point-spread functions. *J. Opt. Soc. Am. A* **19** (2002) 849–857
4. Joshi, N., Szeliski, R., Kriegman, D.J.: Psf estimation using sharp edge prediction. In: *Proc. IEEE Computer Vision and Pattern Recognition Conf.* (2008)
5. Kee, E., Paris, S., Chen, S., Wang, J.: Modeling and removing spatially-varying optical blur. In: *Proc. IEEE Int. Conf. on Computational Photography*. (2011) 1–8
6. Cathey, W.T., Dowski, E.R.: New paradigm for imaging systems. *Applied Optics* **41** (2002) 6080–6092
7. Dorronsoro, C., Guerrero-Colon, J.A., Fuente, M.C., Infante, J.M., Portilla, J.: Low-cost wavefront coding using coma and a denoising-based deconvolution. *Proc. SPIE* **6737** (2007)
8. Chen, J., Yuan, L., Tang, C.K., Quan, L.: Robust dual motion deblurring. In: *Proc. IEEE Computer Vision and Pattern Recognition Conf.* (2008)
9. Zhou, C., Lin, S., Nayar, S.K.: Coded aperture pairs for depth from defocus. In: *Proc. IEEE Int. Conf. on Computer Vision*. (2009)
10. Schechner, Y.Y., Nayar, S.K.: Generalized mosaicing. In: *Proc. IEEE Int. Conf. on Computer Vision*. (2001) 17–24
11. Brown, M., Lowe, D.: Autostitch home page. <http://www.autostitch.net> (2005)
12. Cossairt, O., Nayar, S.K.: Spectral focal sweep: Extended depth of field from chromatic aberrations. In: *Proc. IEEE Int. Conf. on Computational Photography*. (2010)
13. Asada, N., Amano, A., Baba, M.: Photometric calibration of zoom lens systems. In: *Proc. IEEE Int. Conf. on Pattern Recognition*. (1996) 186–190
14. Tang, H., Kutulakos, K.N.: Utilizing optical aberrations for extended-depth-of-field panoramas. Supplementary materials. http://www.cs.toronto.edu/hxtang/projects/edof_panorama/ (2012)
15. Levin, A., Hasinoff, S.W., Green, P., Durand, F., Freeman, W.T.: 4D frequency analysis of computational cameras for depth of field extension. *ACM Transactions on Graphics* **28** (2009)
16. Zhang, Z., Levoy, M.: Wigner distribution and how they relate to the light field. In: *Proc. IEEE Int. Conf. on Computational Photography*. (2009)
17. Ng, R.: Fourier slice photography. *ACM Transactions on Graphics* **24** (2005) 735–744
18. Agarwala, A., Dontcheva, M., Agrawala, M., Drucker, S., Colburn, A., Curless, B., Salesin, D., Cohen, M.: Interactive digital photomontage. *ACM Transactions on Graphics* **23** (2004) 294–302
19. Hasinoff, S., Kutulakos, K.N., Durand, F., Freeman, W.T.: Time-constrained photography. In: *Proc. IEEE Int. Conf. on Computer Vision*. (2009)

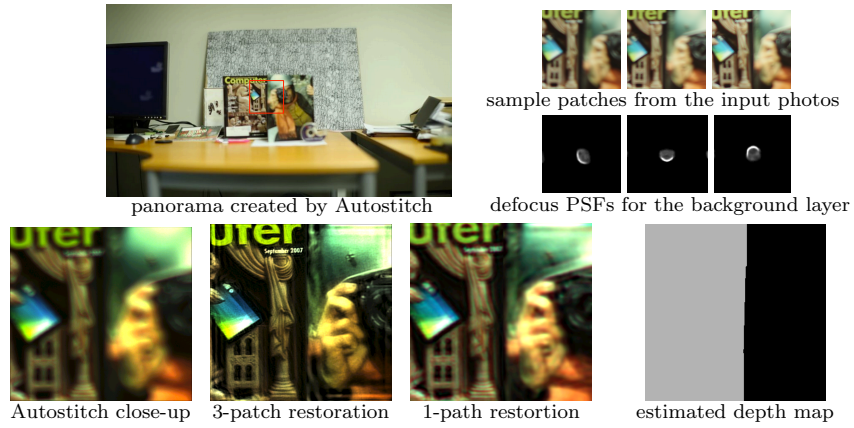


Fig. 8: The “magazine” scene. The lens is focused at $2m$ whereas the patches shown are approximately $1.7m$ away. Note that the lens depth of field at $2m$ is less than $1cm$.



Fig. 9: The “portrait” scene. The lens is focused at infinity with subject at $10m$.

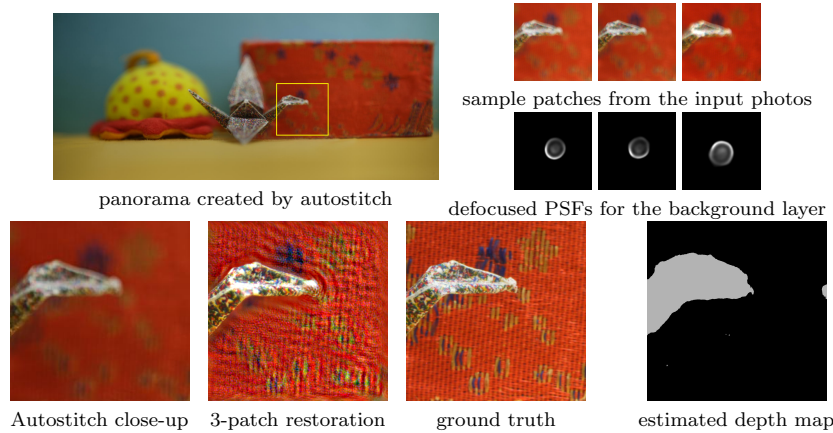


Fig. 10: The “static” scene. The lens is focused at its minimum focusing distance and the foreground region in the patch shown is roughly $3cm$ in front of the in-focus plane.