Modeling and Removing Spatially-Varying Optical Blur

Eric Kee
Dartmouth College
kee@cs.dartmouth.edu

Sylvain Paris
Adobe Systems, Inc.
sparis@adobe.com

Simon Chen
Adobe Systems, Inc.
sichen@adobe.com

Jue Wang
Adobe Systems, Inc.
juewang@adobe.com

Abstract

Photo deblurring has been a major research topic in the past few years. So far, existing methods have focused on removing the blur due to camera shake and object motion. In this paper, we show that the optical system of the camera also generates significant blur, even with professional lenses. We introduce a method to estimate the blur kernel densely over the image and across multiple aperture and zoom settings. Our measures show that the blur kernel can have a non-negligible spread, even with top-of-the-line equipment, and that it varies nontrivially over this domain. In particular, the spatial variations are not radially symmetric and not even left-right symmetric. We develop and compare two models of the optical blur, each of them having its own advantages. We show that our models predict accurate blur kernels that can be used to restore photos. We demonstrate that we can produce images that are more uniformly sharp unlike those produced with spatially-invariant deblurring techniques.

1. Introduction

Many factors can contribute to the undesired blurriness of a photograph. While researchers have well studied blur sources such as camera shake, subject motion, and defocus, and proposed effective solutions to restore the corresponding photos, degradations due to the camera optical system have received little attention. This is particularly unfortunate because optical degradations affect every photograph and cannot be easily removed. This problem is well-known in the photography community as "soft corners" or "coma aberration", and is a discriminating factor between entry-level lenses and professional-grade equipment.

The topic of our study is optical blur. We set up an imaging system in a controlled environment and develop a series of algorithms to extract and evaluate optical blurs that are intrinsic to a particular lens-camera arrangement. Our results show that real optical blur is not only spatially-varying, but also asymmetric. We propose two models to predict the blur kernel (also known as *point spread function* for optical blurs) at any location in the image and for any aperture and focal length, including settings for which we did not make a measurement. This property is key to building a lens profile from few measurements: without it, one would need to fully sample the parameter space of each lens, which is often impractical. Finally, we show that by using the kernels predicted by our models, we can improve the sharpness of the photos captured by corresponding lenses, even when blind deconvolution might fail.

The main contributions presented in this paper include a comprehensive study that reveals unique characteristics of lens blur and two compact optical blur models that enable high quality deblurring results.

1.1. Related Work

Many techniques have been proposed to estimate and remove blur due to camera shake, motion, and defocus, e.g. [1, 2, 4, 6, 8, 12, 14, 16, 18] and references therein. This paper is about blurs that are created in the optics and cannot be significantly improved by focus adjustment. Such blurs have received little attention in comparison. Hsien-Che Lee [10] models a generic optical system, however, our measurements show that there are strong dependencies on the particular lens being used. Sungkil Lee et al. [11] describe optical aberrations from a rendering perspective and note that these effects are often present in real imagery. Gu et al. [5] develop methods to correct dirty or partially occluded optics and Raskar et al. [15] model glare in lenses.

The work that is most related has been done by Joshi et al. [9] to estimate PSFs from edges in the image. They describe how to use a printed pattern to calibrate a camera at a given aperture and focal length and show that they can restore images taken later with the same parameters. We use a similar approach based on a printed pattern. The major improvement brought by our work is that we use our measures to build a parametric model of the spatially-varying optical blur. We show that, with our model, we can restore photos taken with any setting and independently of the image content, which enables the restoration of photos that would be challenging for image-dependent methods, e.g. [9].

2. Modeling Optical Blur

We describe a method to model spatially-varying optical blurs that are intrinsic to a particular lens–camera arrangement. This optical blur varies within an image and depends upon the optical parameters. We model the blur by capturing calibration images from a known lens and camera (Section 2.1), estimating non-parametric blur kernels over local patches of the calibration images (Section 2.2), fitting a parametric blur kernel to each non-parametric kernel (Section 2.3), and modeling how the kernel parameters vary with sensor location and optical settings (Section 2.4).

2.1. Image Capture

We capture images in a controlled environment to isolate optical blur. A planar calibration target is placed in the environment and a camera is mounted on a tripod at a fixed viewing distance. The tripod head is manually adjusted to align the lines of the test chart to the edges of the camera viewfinder. The camera is focused on the center of the test chart using a remote control and images are captured with the internal mirror locked up to eliminate camera vibrations. This configuration is sufficient to calibrate lenses with focal lengths above 17mm and apertures larger than $f/4$, which do not have significant depth-of-field blur in the corners. We capture all images in RAW format at the lowest available ISO setting and estimate blur on RAW color channels prior to color matrixing.

2.2. Non-parametric Kernel Estimation

To model optical blur we begin by computing non-parametric blur kernels that describe the blur in small regions of the test images. Each test image contains a checkerboard test chart with five circles in each square to capture how step-edges of all orientations are blurred. For each square in the test image we align the mathematical definition of the test chart to the local region and synthesize a sharp square (Section 2.2.1). We use the test image and synthesized sharp square to estimate a non-parametric kernel (Section 2.2.2). This process is summarized in Figure 1.

2.2.1 Test Chart Alignment

The test chart is aligned to individual squares in the test image before estimating the corresponding blur kernel. The corners of the square in the test image are localized and a bootstrap projective homography H is computed that aligns the test chart to the square [7]. The homography is used to rasterize a synthetic image of the square from its mathematical definition: chart edges are anti-aliased by computing the fraction of the pixel that is filled and the white/black point are set to the histogram peaks of the square in the test image. This shading is effective because it is applied locally. The homography is then iteratively refined. In the i^{th} iteration, a coarse-to-fine differential registration technique is used to compute a projective homography H_i that aligns the synthetic square to the observed image. The homography is then updated, $H \leftarrow H_i H$ and the synthetic image is re-rasterized. This iteration ends when H_i is close to an identity. The resulting homography gives sub-pixel alignment between the test chart and blurry square, Figure 1.

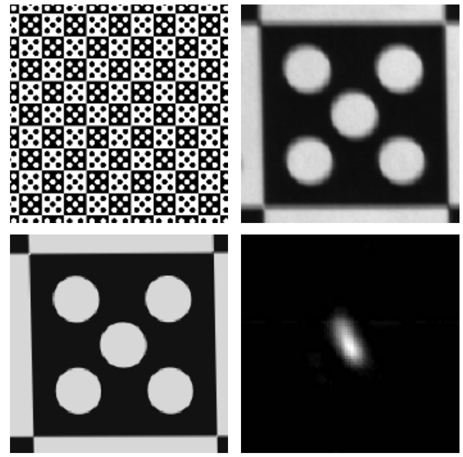

Figure 1. Blur estimation. Top left, the test chart; top right, a test image with optical blur; bottom left, the synthesized, aligned, test chart; bottom right, the 3x super-resolved non-parametric kernel.

2.2.2 Kernel Estimation

We estimate non-parametric blur kernels for each square by synthesizing a sharp square from the aligned test chart. The blur kernel can be computed by using conjugate gradient descent to solve the least squares system $Ak = b$, where k is the kernel, A is a Toeplitz matrix that encodes the convolution of the sharp square with the kernel, and b is the blurry square. This optimization can be performed efficiently in the Fourier domain without explicitly constructing A. Although this method allows negative kernel values, in practice these are small and easily removed by thresholding and re-normalizing the kernel.

Because optical blurs are sometimes small, in practice we super-resolve the blur kernel. The homography H, which is known to sub-pixel accuracy, is used to synthesize a high-resolution test chart and the linear system becomes $WA_r k_r = WUb$, where A_r and k_r encode the high-resolution test chart and kernel. Matrix U up-samples b and W is a weight matrix that assigns zero-weight to interpolated pixels. By formulating this problem with U and W, matrix A_r does not need to be constructed and the convolutions can be performed in the Fourier domain. This computation is fast compared to non-negative least squares, as in [9], and a smoothness regularization term was not necessary. To account for the interpolant when deblurring, in practice we estimate kernels with a uniform weight matrix $W = I$. This additional smoothness does not distort the kernel when observed at the original resolution. In this work we super-resolve kernels at 3x image resolution.

2.3. Parametric Kernel Fitting

Non-parametric kernels can describe complex blurs but their high dimensionality masks the relationship between the kernel shape and the optical parameters. We use a 2–D Gaussian distribution to reduce the dimensionality and model the kernel shape. Because non-parametric kernels can be noisy, we use a robust method to fit the 2–D Gaussian. The non-parametric kernel is thresholded, isolated regions are labelled, and the maximum likelihood (ML) estimator is used to fit a 2–D Gaussian to the central region. The ML Gaussian then iteratively refined by using the Levenberg-Marquardt algorithm to estimate the Gaussian parameters that minimize the SSD error between the non-parametric kernel and the synthesized 2–D Gaussian.

To quantify the impact of the Gaussian approximation and validate the robust fitting method, we compared images that were deconvolved with non-parameteric kernels, ML Gaussian kernels, and robust-fit Gaussian kernels. Specifically, we used two images, one natural and one synthetic, and called them sharp. We blurred both sharp images with 660 non-parametric kernels that were estimated from a test lens. This produced two sets of 660 blurry images. We deconvolved each of the blurry images with the requisite non-parametric kernel, robust-fit Gaussian kernel, and ML Gaussian kernel.

We measured the visual quality of the deconvolutions by computing the mean Structural Similarity Index (SSIM) [17] between the deconvolved and sharp images. If the ML Gaussian kernels produce larger error than the robust-fit Gaussians, the SSIM index of the robust-fit Gaussian is greater. However, the difference between the two SSIM indices is also greater when kernels are large because deconvolution error increases with kernel size, even when using the ground-truth kernel [14]. Therefore, we compare SSIM indices and account for kernel size by comparing a ratio, as in [14],

$$\text{error ratio} = (\text{SSIM}_{\text{non}} + 2)/(\text{SSIM}_{\text{gau}} + 2), \qquad (1)$$

where SSIM_{non} is the SSIM given by the non-parametric kernel and SSIM_{gau} is the SSIM given by a Gaussian kernel. (The $+2$ is added to shift the SSIM into $[1,\ 2]$). This error ratio is greater than one when deconvolution by the Gaussian kernel produces worse results than the non-parametric kernel and is equal to one when the Gaussian kernel produces an identical result.

Figure 2 (bottom) shows the cumulative distribution of the errors for the ML Gaussians (dashed lines) and robust-fit Gaussians (solid lines). The robust-fit Gaussians produce lower deconvolution error: 99% of the errors fell below 1.01 and 1.02 for the natural and synthetic images, respectively. For the ML Gaussians, 99% of the errors fell below 1.15 and 1.26.

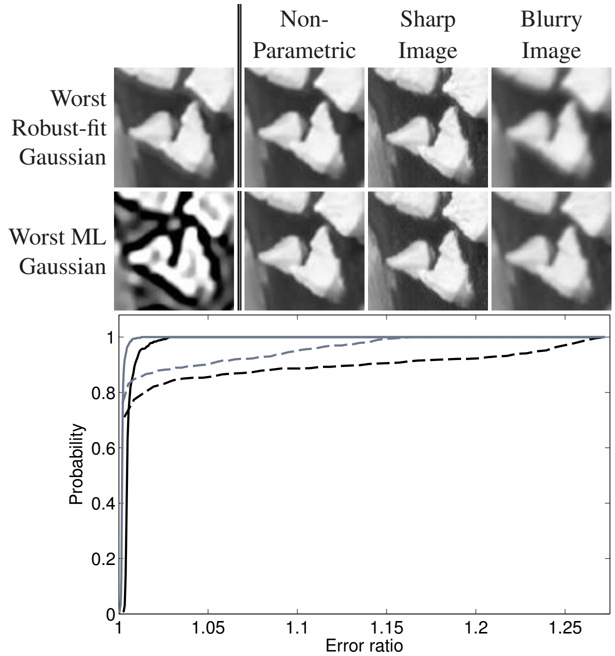

Figure 2. Gaussian approximation error using robust and maximum likelihood Gaussians. Top, deconvolution with the Gaussian kernels that gave the largest (worst) SSIM error ratio of all 660 kernels. Bottom, cumulative distributions of error ratios. Dashed lines, using the ML Gaussian; solid lines, the robust Gaussian. Light/dark lines show error in the natural/synthetic image.

Figure 2 (top) shows the deconvolved natural images that gave the largest (worst) SSIM error ratio for both types of Gaussian kernels. The worst robust-fit Gaussian kernel produces a deconvolution result that is very similar to deconvolution by the non-parametric kernel and improves the blurry image. In contrast, the worst ML Gaussian produces a deconvolution result with dramatic artifacts that are not present when using the non-parametric kernel. The worst ML Gaussian result (SSIM error ratio 1.17), and worst robust-fit Gaussian result (SSIM error ratio 1.01), give intuition for the range of visual errors along the x-axis of the cumulative distribution (bottom). This demonstrates that the robust-fit Gaussians produce visually small errors and are a good approximation to the optical blurs in the test lens.

2.4. Kernel Variation Models

Optical blur depends upon multiple factors including spatial location on the sensor x, y, focal length f, aperture a, and color channel. Images may contain significant, asymmetric, optical blurs, particularly in the corners, Figure 3. A calibration image could be used to estimate and correct such blurs for photos taken with the same lens settings; however, in practice it is difficult or impossible to calibrate all possible settings. To overcome this problem we developed two models, each with different strengths, to predict optical blur at novel settings from a sparse set of calibration images. Specifically, we compare two models to describe how the Gaussian covariance parameters

$$\Sigma = \begin{bmatrix} C_{xx}^2 & C_{xy} \\ C_{xy} & C_{yy}^2 \end{bmatrix} \tag{2}$$

vary. The mean is assumed to be zero. Both models comprise three independent polynomials that predict the three degrees of freedom in Σ, C_{xx}, C_{yy}, and the correlation $C_{or} = C_{xy}/C_{xx}C_{yy}$. We model the blur in each color channel separately and the remaining discussion addresses a single color channel.

The first model is a polynomial $G(x, y, f, a)$, where x, y are the spatial location of the kernel and f, a are the focal length and aperture at which the image was captured. Specifically, this global model may be any $G_{\alpha,\beta}$ that contains all polynomial terms in x, y, f, a up to order $\max(\alpha, \beta)$, such that the order of terms that contain x or y is at most α and the order of terms that contain f and a is at most β. For example, $G_{3,1}$ contains all third order polynomial terms in x, y, f, a, excepting terms such as a^2x. Intuitively, α and β limit the complexity of the individual polynomial predictors that compose the model. Parameters α and β were selected to control x, y and f, a because our experiments show that optical blur is complex in x, y yet simple in f, a (see Section 3).

The second model is a polynomial $L(x, y, f, a)$ that describes the blur in a local x, y region. The local model has the same form as the global: any $L_{\alpha,\beta}$ that contains all polynomial terms in x, y, f, a up to order $\max(\alpha, \beta)$, where α and β control the complexity of the model in x, y and f, a. The motivation behind this local model is illustrated in Figure 3. The relationship between optical blur and x, y may be complex and require large α for $G_{\alpha,\beta}$; however, small α may be reasonable for a local region. Because we can easily collect dense blur samples by decreasing the test chart scale, the local model takes a more data-driven approach and fits $L_{\alpha,\beta}$ to local regions in which α is small.

The radius of the local model and the test chart resolution are chosen to match the complexity of the spatial variations. For stability, the radius should be at least 2x the width of the largest imaged squares. A radius of 3x worked well. Imaged squares were between 2% and 6% of image width and we sampled the blur at alternating squares.

2.4.1 Model Selection

To fit the blur models, images are captured at multiple focal length and aperture combinations. The sampling resolution of f and a is estimated by cross-validation over α and β but a starting point is needed for data collection. We began by collecting a set of images at a fixed focal length and varying aperture. For each image, we computed Gaussian kernels and plotted C_{xx}, C_{yy}, and C_{or} at each x, y location as a function of aperture. We repeated this process for a fixed aperture and varying focal length and used these aperture and focal length sweep plots to select an initial sampling resolution in f, a according to the rate at which the Gaussian parameters vary, Figure 4 (details in Section 3).

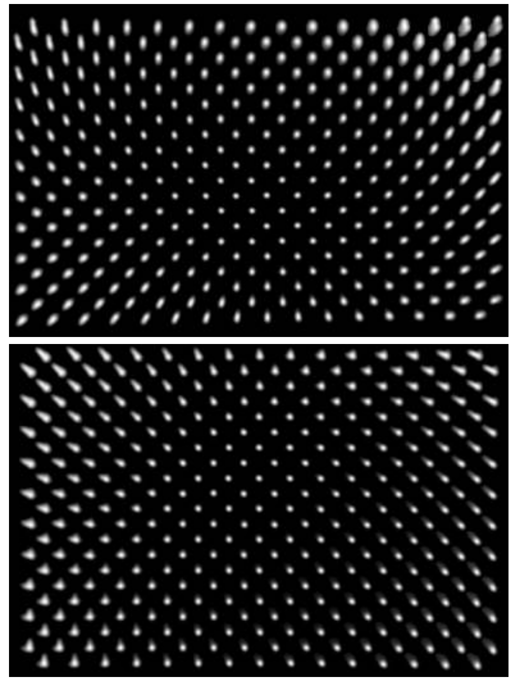

Figure 3. Non-parametric kernels from two Canon 17-40mm $f/4$ lenses on a Canon 1D Mark III (at 40mm $f/5$). Kernels are super-resolved at 3x and enlarged for display. The maximum blur of the Gaussian approximations has standard deviation of 3 pixels; the minimum is 1 pixel standard deviation.

The complexity of the global and local models is determined by cross-validation over α and β. A cross-validation dataset of $N_f \times N_a$ images is captured, where N_f and N_a are the number of samples in the f and a dimensions. The sampling resolution constrains the complexity of $G_{\alpha,\beta}$ and $L_{\alpha,\beta}$ to $\beta < \min(N_f, N_a)$. In this work, we consider locally-linear models $L_{1,\beta}$ and perform cross-validation of $L_{1,\beta}$ for β alone.

The values of α,β for the global model are computed in two cross-validation stages. Let $N = \min(N_f, N_a)$. First we plot the mean prediction error of $G_{\alpha,N-1}$ against α and select an optimal value α_{opt}. In the second stage, we repeat this analysis for $G_{\alpha_{\text{opt}},\beta}$, $\beta < N$, and select β.

3. Model Complexity and Comparison

We estimated the focal length and aperture sampling resolution using a Canon 24–105mm $f/4$ (MSRP $\$1,249$) on a Canon 1D Mark III, (MSRP $\$3,999$). Figure 4 shows the aperture and focal length sweep plots for C_{xx}. Each line represents C_{xx} measured at a fixed sensor location. We performed the aperture sweep (left) by capturing images at each aperture, $f/4$ to $f/10$. The focal length sweep (right) was performed across the focal length range. Sharp variations in the focal length sweep are noise caused by manually changing the focal length. Consequently, we subsampled the aperture to $N_a = 5$ settings, $f/\{4,\ 5,\ 6.3,\ 8,\ 10\}$, and $N_f = 5$ focal lengths, $\{24,\ 44,\ 65,\ 85,\ 105\}$mm.

We computed the complexity of the models by cross-validation using the Canon 24–105mm $f/4$ and a second lens, a Canon 17–40mm $f/4$ (MSRP $\$840$). Similarly to the 24–105mm lens, we sampled the 17–40mm at 5 settings in focal length and aperture. We collected 50 images from each lens, two images at each setting[1]. We computed the Gaussian kernels for each image and used holdout-10 cross validation to compute prediction error.

Figure 5 shows cross-validation error for the global model: left, error for α; right, β. The local-linear plots for β are similar to those in Figure 5 (right)[2]. We compute the error as a percentage of the range of blurs on each lens and define this range to be the width of the 99% confidence interval for each Gaussian parameter. Not shown in Figure 5 are the error distributions at each α, β. On the 17–40mm lens, 95% of the error is below 20% for $\alpha > 6$; on the 24–105mm lens, $\alpha > 5$.

Based upon the mean and distributions of the cross-validation error, we selected models $G_{8,3}$ and $L_{1,3}$.

We compared the global and locally-linear models using a test set of 16 images that we captured at novel settings of the Canon 17–40mm $f/4$ lens: $\{19.8,\ 25.6,\ 31.37,\ 37.1\}$mm and $f/\{4.5,\ 5.6,\ 7.1,\ 9.0\}$. We estimated the Gaussian kernels in each test image, fit models $G_{8,3}$ and $L_{1,3}$ to the cross-validation dataset, and computed prediction error. Figure 6 shows the cumulative distribution of error for $G_{8,3}$ (dark solid lines) and $L_{1,3}$ (dark dashed lines). Top left, C_{xx}; right, C_{yy}; bottom, C_{or}. For all parameters, 95% of errors were below 10%. Notably, models $G_{8,1}$ and $L_{1,1}$ (light lines) are also competitive.

To summarize, Figure 4 shows that optical blur varies slowly with f, a and these dimensions may be subsampled. Figure 5 shows that optical blur is complex in x, y, simple in f, a, and we select models $G_{8,3}$ and $L_{1,3}$. Figure 6 shows that $G_{8,3}$ and $L_{1,3}$ are equally good models for the test lens.

[1] Samples from the cross-validation set are shown in Figures 1 and 7.

[2] At $L_{1,3}$, for the Canon 24–105mm, mean cross-validation error was 3.6%, 3.7%, and 3.4% (C_{xx} C_{yy} and C_{or}); for the Canon 17–40mm it was 4.0%, 3.5%, and 3.7%.

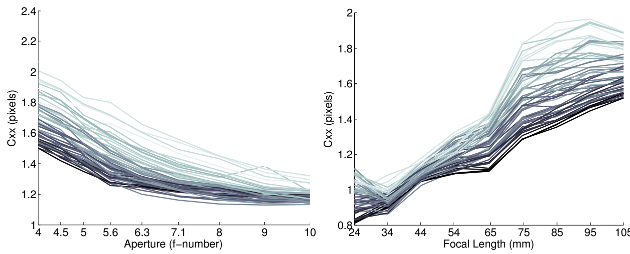

Figure 4. Aperture and focal length sweeps for C_{xx} using a Canon 24–105mm $f/4$ lens on a Canon 1D Mark III. Each line represents the C_{xx} term at a fixed sensor location. The x-axis labels mark the sampled parameter values. Left, C_{xx} at 105mm from $f/4$ to $f/10$; right, C_{xx} at $f/4$ from 24–105mm.

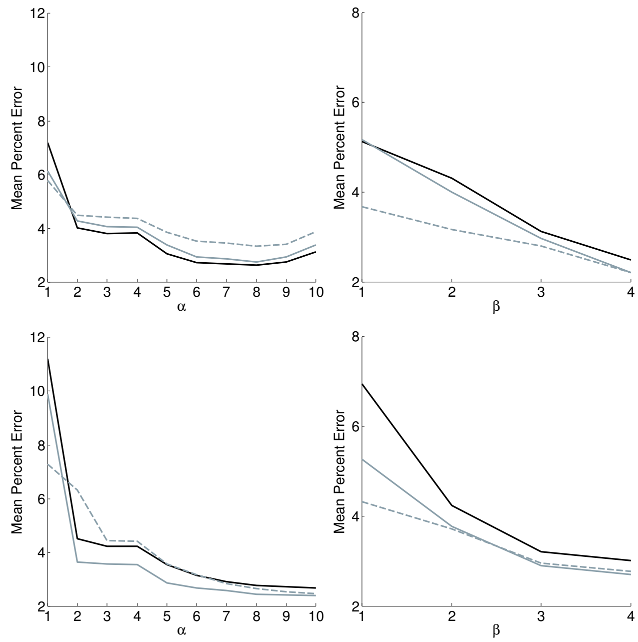

Figure 5. Mean cross-validation error for the global model using two lenses. Top row, Canon 24–105mm $f/4$; bottom row, Canon 17–40mm $f/4$. Dark and light solid lines represent C_{xx} and C_{yy} error; dashed lines, C_{or} error. Left, mean percent error for models $G_{\alpha,4}$. Right, mean percent error for models $G_{8,\beta}$.

4. Results

To give intuition for the visual impact of prediction error, we selected Gaussian kernels with an average error of 1%, 10%, and maximum error under the $L_{1,3}$ model. For each kernel, we deconvolved the requisite square in the test image using the predicted Gaussian kernel, the robust-fit Gaussian kernel, and the non-parametric kernel, Figure 7. We used the non-blind deconvolution algorithm described in [12] (code: [13]). At 1% error (top row), the predicted and robust-fit kernels produce nearly identical results. At 10% error (middle row), the predicted kernel produces a

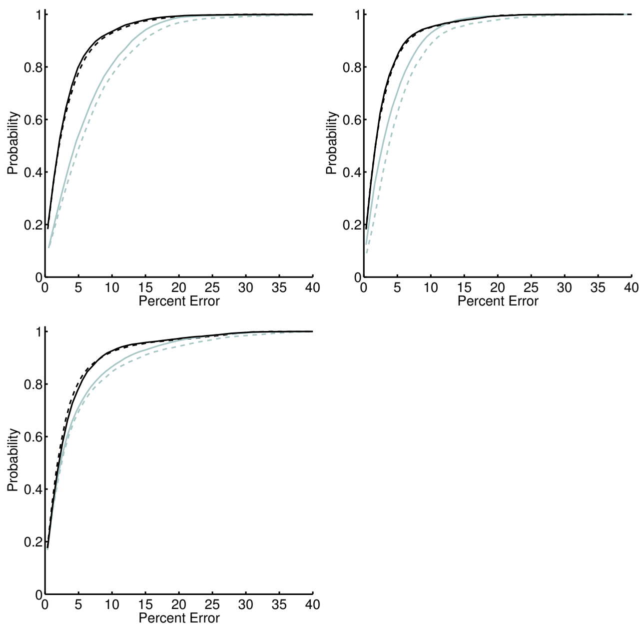

Figure 6. Cumulative distribution of test error when predicting blurs at novel lens settings. Top left, C_{xx}, right C_{yy}, bottom C_{or}. Solid lines denote global models; dashed lines denote locally-linear models. Dark lines denote $G_{8,3}$ and $L_{1,3}$ models; light lines denote $G_{8,1}$ and $L_{1,1}$ models.

sharper result than the robust-fit kernel, demonstrating that the model reduces noise in the individual samples. At 40% error (bottom row), both the predicted and robust-fit kernels produce blurrier results than the non-parametric kernel.

To test deblurring outside of the lab, we captured two images using the Canon 17–40mm $f/4$ on a Canon 1D Mark III. The first is a dominantly planar indoor scene, Figure 8 (top). We placed the camera on a flat surface, supported the body manually, used mirror-up mode to reduce vibration, and focused on a location near 11 on the center clock dial. Spatially-varying blur can be seen at locations 1 and 2, Figure 8 (middle and bottom). We deconvolved both locations using the Gaussian kernel predicted by $L_{1,3}$ (figure label b) and a non-parametric kernel (figure label a) estimated from a calibration image. The predicted kernel produces a result that is very similar to the non-parametric kernel.

For additional comparison, we deblurred the indoor image with the optical deblurring software, DxO [3]. We adjusted the parameters for the best output, Figure 11. Compare the DxO result to the result when using a ground truth kernel taken from a calibration image and when using the predicted Gaussian. The DxO result is blurrier, possibly because one model is used for all Canon 17–40mm $f/4$ lenses.

We captured an outdoor image with the same setting as the indoor image, 35mm $f/4$, Figure 9 (top). Spatially varying blur can be seen at locations 1 and 2. The image is sharper when deconvolving with both the $L_{1,3}$ Gaussian and the non-parametric kernel (labels b and a).

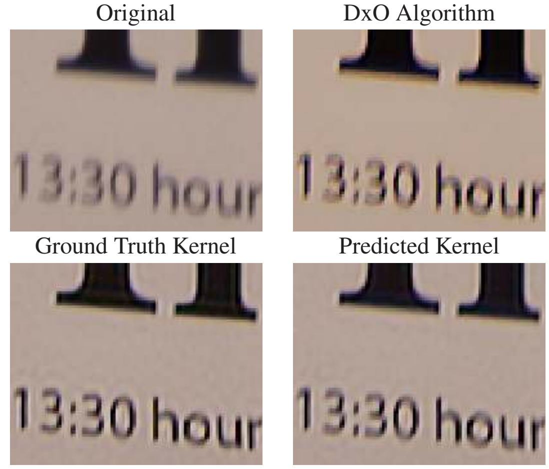

Figure 11. Comparison to DxO software [3]. A non-parametric kernel from a calibration image is used as ground truth. The $L_{1,3}$-predicted Gaussian kernel produces a sharper solution than DxO.

We also tested a consumer-grade system, a Canon Rebel T2i with a 18–55mm $f/3.5$–5.6 lens (MSRP \$899 combined). Deblurring results at 18mm $f/3.5$ are shown in Figure 10. Blurs for the Rebel are spatially-varying and smaller at 18mm $f/3.5$ than the Canon 17–40mm (relative to the much larger image size of the T2i).

The full resolution images for Figures 8–11 are available at: `http://www.juew.org/lensblur/materials.zip`

Finally, we tested two additional lenses, a Nikkor 24–120mm $f/3.5-5.6$ (MSRP \$669) on a Nikon D3 (MSRP \$4,999), and a duplicate Canon 17–40mm $f/4$. Prediction error was low for the Nikkor: 95% of predictions were within 10% error. We used the duplicate Canon lens to test if optical blur varies across lenses of the same make/model. Using the first Canon model to predict blurs in the duplicate gave large error: 95% of prediction errors were below 51%, 46%, and 74% error[3]. The blur in the duplicate Canon is quantitatively and qualitatively different, Figure 3.

5. Discussion

The results show that our optical models provide accurate kernels for image restoration and successfully interpolate data between measurement points. Practically, the size of the models—global $\approx$ 12kb, local $\approx$ 200kb—is sufficiently small to be included in photo editing packages. Furthermore, the low-order relationship between optical blur and focal length/aperture allows both models to be fit with few calibration images and the Gaussian parameters can be efficiently estimated directly from an image.

An intriguing point is that the kernels that we measured differ from Joshi's [9], which are more disc-shaped. One

[3] Prediction error is defined in Section 3.

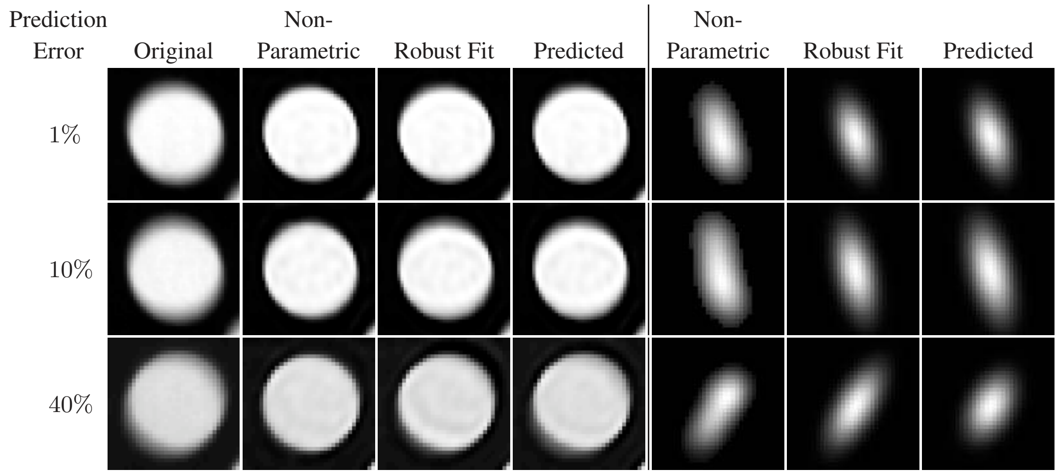

Figure 7. Deconvolution at varying prediction error. First column, the mean prediction error of the sample's C_{xx}, C_{yy}, C_{or}. Left half, deconvolution using each kernel type; right half, the kernels.

explanation is that Joshi and colleagues studied consumer-grade lenses that may suffer from front- or back-focusing. In this case, their measures would be dominated by defocus blur that corresponds to the image of the aperture possibly truncated by the lens barrel. Our pro-grade lenses are likely to focus accurately and greatly reduce defocus blur. Our observations are mostly due to imperfections in the glass used to build the lens. This difference is also consistent with the fact that Joshi's kernels are symmetric because they depend on the shape of the barrel, whereas ours are not because they are due to inaccuracies in the glass. We believe that it is important to model the optical inaccuracies because, as we have shown, these visibly degrade image sharpness.

References

[1] S. Cho and S. Lee. Fast motion deblurring. In *SIGGRAPH Asia '09: ACM SIGGRAPH Asia 2009 papers*, pages 1–8, New York, NY, USA, 2009. ACM.

[2] T. S. Cho, A. Levin, F. Durand, and W. T. Freeman. Motion blur removal with orthogonal parabolic exposures. In *IEEE International Conference in Computational Photography (ICCP)*, 2010.

[3] DxO. Dxo optics. In *http://www.dxo.com*, 2010.

[4] R. Fergus, B. Singh, A. Hertzmann, S. T. Roweis, and W. T. Freeman. Removing camera shake from a single photograph. *ACM Trans. Graph.*, 25(3):787–794, 2006.

[5] J. Gu, R. Ramamoorthi, P. Belhumeur, and S. Nayar. Removing Image Artifacts Due to Dirty Camera Lenses and Thin Occluders. *ACM Trans. Graph.*, Dec 2009.

[6] A. Gupta, L. Joshi, N.and Zitnick, M. Cohen, and B. Curless. Single image deblurring using motion density functions. In *Proceedings of European Conference on Computer Vision*, 2010.

[7] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521540518, second edition, 2004.

[8] N. Joshi, S. B. Kang, C. L. Zitnick, and R. Szeliski. Image deblurring using inertial measurement sensors. *ACM Trans. Graph.*, 29(4):1–9, 2010.

[9] N. Joshi, R. Szeliski, and D. Kriegman. PSF estimation using sharp edge prediction. In *IEEE Conference on Computer Vision and Pattern Recognition*, June 2008.

[10] H.-C. Lee. Review of image-blur models in a photographic system using the principles of optics. *Optical Engineering*, 29(05), 1990.

[11] S. Lee, E. Eisemann, and H.-P. Seidel. Real-Time Lens Blur Effects and Focus Control. *ACM Trans. Graph.*, 29(4):65:1–7, 2010.

[12] A. Levin, R. Fergus, F. Durand, and W. T. Freeman. Image and depth from a conventional camera with a coded aperture. *ACM Trans. Graph.*, 26, July 2007.

[13] A. Levin, R. Fergus, F. Durand, and W. T. Freeman. Image and depth from a conventional camera with a coded aperture. In *http://groups.csail.mit.edu/graphics/CodedAperture*, 2007.

[14] A. Levin, Y. Weiss, F. Durand, and W. Freeman. Understanding and evaluating blind deconvolution algorithms. *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*, 0:1964–1971, 2009.

[15] R. Raskar, A. Agrawal, C. A. Wilson, and A. Veeraraghavan. Glare aware photography: 4d ray sampling for reducing glare effects of camera lenses. *ACM Trans. Graph.*, 27:56:1–56:10, August 2008.

[16] Q. Shan, J. Jia, and A. Agarwala. High-quality motion deblurring from a single image. *ACM Trans. Graph.*, 27(3):1–10, 2008.

[17] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.

[18] O. Whyte, J. Sivic, A. Zisserman, and J. Ponce. Non-uniform deblurring for shaken images. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2010.

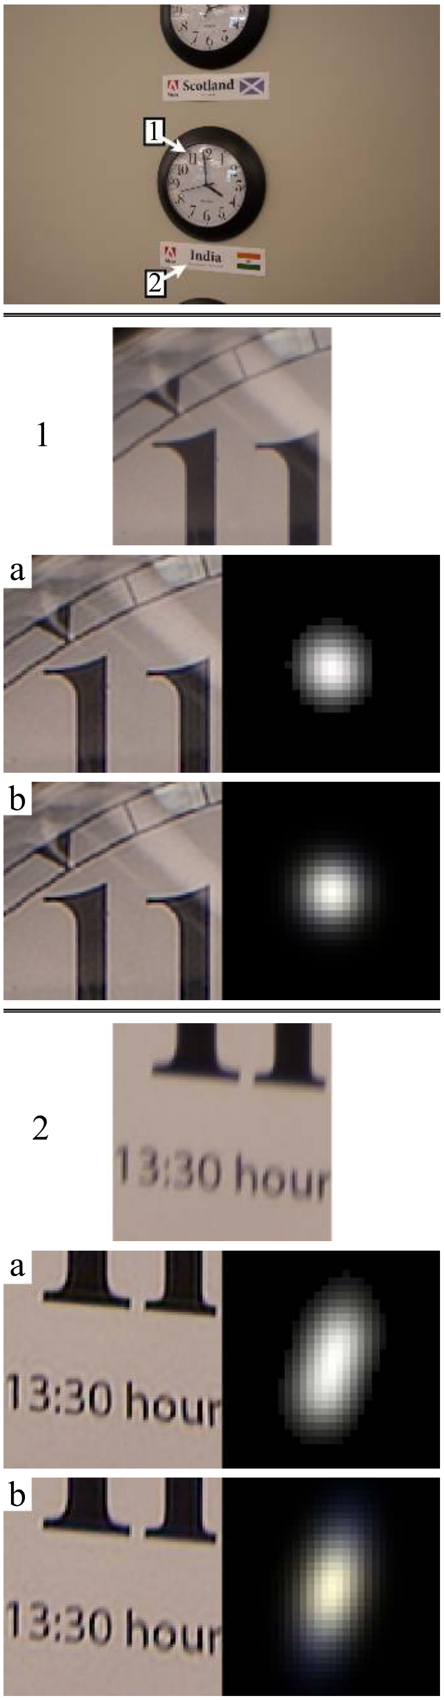

Figure 8. Image taken with a Canon 1D Mark III, at 35mm $f/4.5$. Images a, b are deblurred with non-parametric and Gaussian kernels.

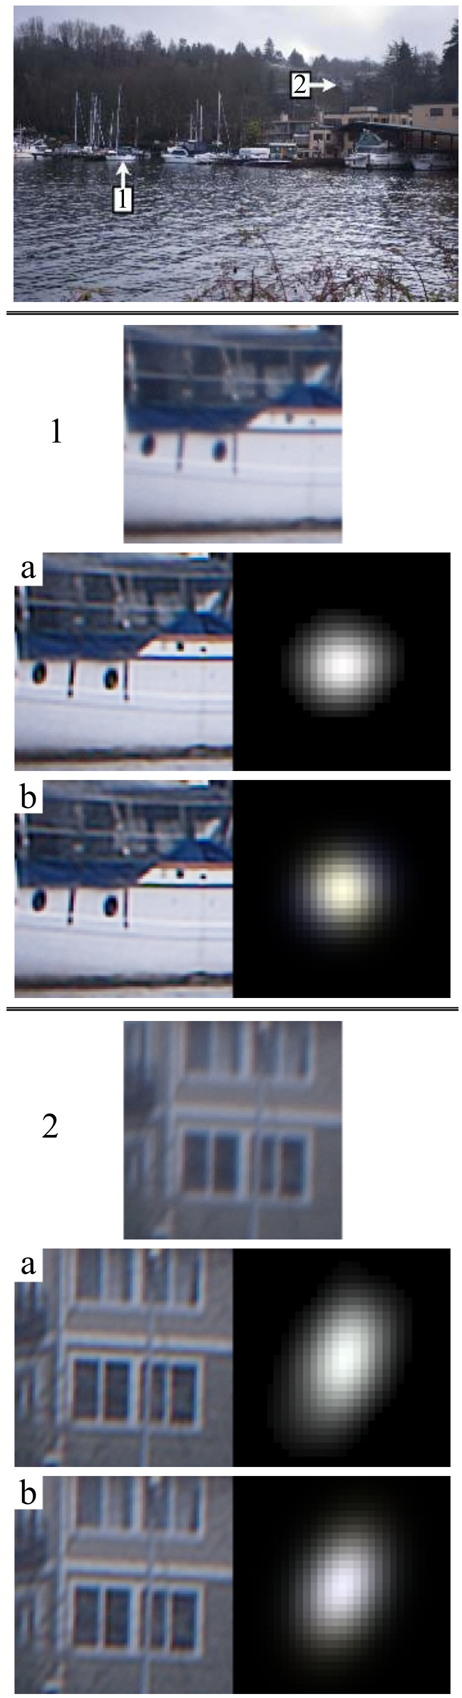

Figure 9. Image taken with a Canon 1D Mark III, at 35mm $f/4.5$. Images a, b are deblurred with non-parametric and Gaussian kernels.

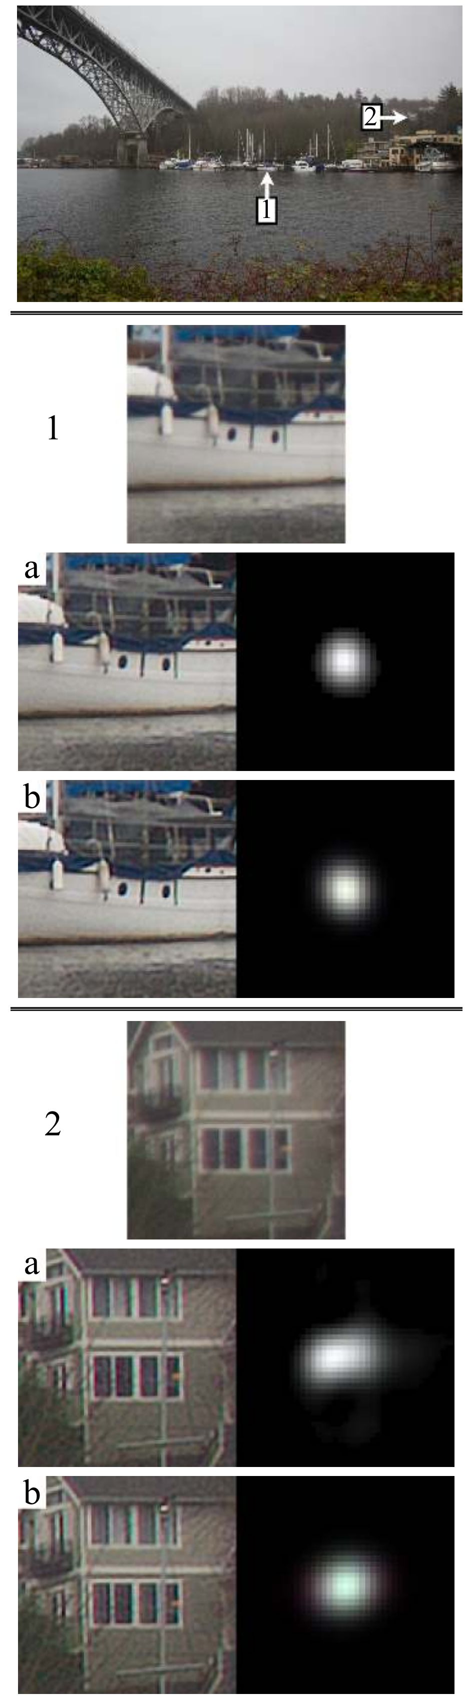

Figure 10. Image taken with a Canon Rebel T2i at 18mm $f/3.5$. Images a, b are deblurred with non-parametric and Gaussian kernels.