What are Good Apertures for Defocus Deblurring?

Changyin Zhou, Shree Nayar *†

Abstract

In recent years, with camera pixels shrinking in size, images are more likely to include defocused regions. In order to recover scene details from defocused regions, deblurring techniques must be applied. It is well known that the quality of a deblurred image is closely related to the defocus kernel, which is determined by the pattern of the aperture. The design of aperture patterns has been studied for decades in several fields, including optics, astronomy, computer vision, and computer graphics. However, previous attempts at designing apertures have been based on intuitive criteria related to the shape of the power spectrum of the aperture pattern. In this paper, we present a comprehensive framework for evaluating an aperture pattern based on the quality of deblurring. Our criterion explicitly accounts for the effects of image noise and the statistics of natural images. Based on our criterion, we have developed a genetic algorithm that converges very quickly to near-optimal aperture patterns. We have conducted extensive simulations and experiments to compare our apertures with previously proposed ones.

1. Introduction

Since the 1990s, the spatial resolution of image detectors has been increasing at a rapid pace. This trend is being driven by advances in silicon technology that enable the fabrication of smaller pixels. For a given optical setting, smaller pixels result in a smaller depth of field (DOF). Interestingly, smaller pixels need more light to maintain signal-to-noise ratio (SNR) and hence require the use of wider apertures, which causes further reduction in DOF. The end result is that, with increase in resolution, images are more inclined to include large defocused regions, where scene details are blurred out. The only way to recover these details is by using deblurring techniques. For these reasons, image defocus deblurring has recently resurfaced as an active area of research.

It is well-known that out-of-focus blurring can be formulated as a convolution of the perfectly focused image with a kernel that is determined by an aperture pattern; defocus deblurring is achieved by deconvolution with this kernel. The main problem with defocus deblurring is that the higher frequencies of the signal are attenuated during image formation and consequently deconvolution amplifies image noise. For any given frequency in the Fourier domain, the lower the power the defocus kernel has, the greater the amplification of image noise. In the case of a conventional circular aperture, the defocus kernel is known to not only severely attenuate high frequencies but also have zero-crossings in frequency domain. Over the past 50 years, numerous aperture designs have been proposed to preserve high frequency information (e.g. [1] [2]). In recent years, new coded apertures have been proposed for defocus deblurring [3]. These works have evaluated and optimized aperture patterns based on intuitive criteria related to the shape of their power spectra, such as maximizing the minimum value of the spectrum [3]. Although such criteria have helped to find better aperture patterns, they do not explicitly account for the effects of image noise and image structure in the context of defocus deblurring.

In this paper, the "goodness" of an aperture pattern is evaluated based on the quality of deblurring, rather than on any particular characteristic of the aperture pattern's power spectrum. In our method, the spectrum of an aperture pattern is assessed together with the level of image noise and the expected spectrum of an image. For the image spectrum we use the well-known $1/f$ law [4][5][6] as an image prior. Despite the fact that our apertures are optimized using this specific prior, we have found that they produce high quality deblurrings for a wide variety of real-world images.

Even though our evaluation criterion is concise, finding the optimal pattern is still a challenging problem. For a binary pattern of resolution $N \times N$, the number of possible solutions is $2^{N \times N}$. This makes finding the optimal pattern intractable. To solve this optimization problem, we use a genetic algorithm [7] in which the pattern is represented by a gene sequence that evolves via selection, turnover, and mutation. Because of the simplicity of our pattern evaluation criterion and the efficiency of the proposed genetic algorithm, for a 13×13 pattern, the optimization converges to a near-optimal solution in about 20 minutes on a 4GHz PC.

*Changyin Zhou and Shree Nayar are with the Department of Computer Science, Columbia University, New York, NY, 10027

†Email: {changyin, nayar}@cs.columbia.edu

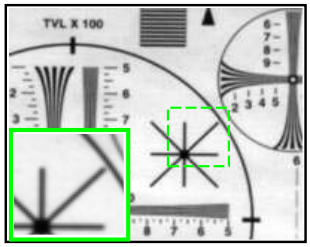

(a) Focused Image (Ground Truth)

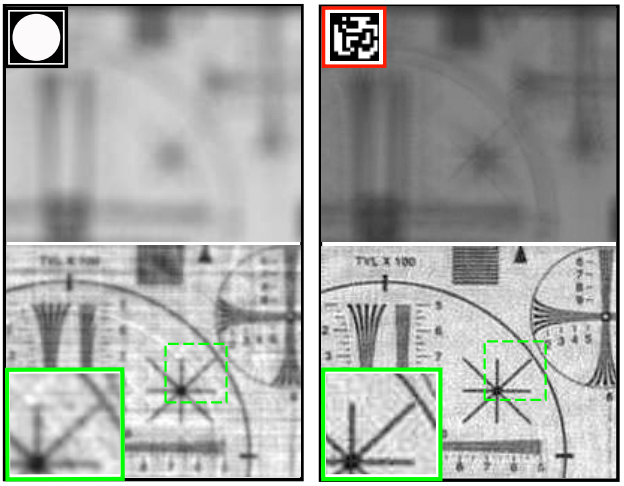

(b) Circular Aperture (c) Proposed Aperture

Figure 1. Comparison of deblurring results obtained using a circular aperture and one of our optimized apertures. (a) A focused image of a CZP resolution chart. (b) Severely defocused image captured using a circular aperture (top) and the result of deblurring (bottom). (c) Image captured using our optimized aperture (top) and the result of deblurring (bottom). The apertures used are shown in the top-left corners of the captured images. Both the captured images were taken under identical focus and exposure settings (hence the darker captured image in (c)).

To experimentally verify our optimized patterns, we printed several aperture patterns as high resolution (1 micron) photomasks and inserted them into Canon EF $50mm$, $f/1.8$ lenses. These lenses were attached to a Canon EOS 20D camera and used to capture images of a wide variety of scenes. For example, Figure 1 compares the deblurring results for a CZP resolution chart obtained with a circular aperture and our optimized aperture. We can see that although the captured image is highly defocused, most details are recovered when the optimized aperture is used. In the case of the conventional circular aperture, however, a significant amount of information is lost – the deblurring result is very noisy, is lacking in high frequencies, and includes many artifacts.

Given an aperture pattern, we still need the scene depth to determine the size of the kernel to deblur with. In this paper, we focus on the problem of how to best preserve information during out-of-focus blurring by choosing proper aperture patterns, and have assumed that scene depth is provided either manually or by a depth estimation method. When depth information is not available, users can try different scene depths until scene details are best recovered. The quality of these recovered details, such as car license numbers, telephone numbers and human faces, can be critical in a variety of imaging applications.

2. Related Work

In the early 1960s, coded aperture techniques were introduced in the field of high energy astronomy as a novel way of addressing the SNR issues related to lensless imaging of x-ray and γ-ray sources [8]. In subsequent decades, many different aperture patterns were proposed, including the popular modified uniformly redundant array (MURA) [9]. Unfortunately, the coded apertures designed for lensless imaging are not optimal to use with lenses for defocus deblurring, as observed in [3].

Also in the 1960s, researchers in the field of optics began developing unconventional apertures to increase DOF as well as capture high frequencies with less attenuation [1][2]. These apertures were usually chosen based on simple intuitions and then analyzed in terms of their optical transfer functions. A different set of approaches use a 3D phase plate at the aperture plane [10],[11], or a moving image detector [12], to extend DOF. The goal of these approaches is to make the blur kernel depth-invariant rather than optimal for defocus deblurring.

It is only in the last few years that the design of apertures for defocus deblurring was posed as an optimization problem. In particular, Veeraraghavan et al. [3] used gradient descent search to improve the MURA pattern [9] and then binarized the resulting pattern. Due to the large search space associated with the optimization, they restricted themselves to binary patterns with 7×7 cells. The criterion used in [3] maximizes the minimum of the power spectrum of the aperture pattern.

In another related work by Levin et al. [13], the aperture pattern is optimized for the recovery of depth from defocus, a different problem from the one we address. Since they also use their optimized pattern for defocus deblurring in their experiments, we include their pattern in our comparisons. However, to be fair, it should be noted that their pattern was not designed for defocus deblurring. It is worth mentioning that patterned apertures have also been used in other imaging applications [14] [15] [16] [17].

3. Criterion for Aperture Quality

3.1. Formulating Defocus Deblurring

For a simple fronto-planar object, its out-of-focus image can be expressed as:

$$f = f_0 \otimes k + \eta\,, \tag{1}$$

where, f_0 is the focused image, k is the point spread function (PSF) determined by the aperture pattern and the degree of defocus, and η is the image noise which is assumed to be Gaussian white noise $N(0, \sigma^2)$. In frequency domain, we

have

$$F = F_0 \cdot K + \zeta, \tag{2}$$

where, F_0, K, and ζ are the discrete Fourier transforms of f_0, k, and η, respectively.

Given a defocused image F and known PSF K, the problem of defocus deblurring is to estimate the focused image F_0 by solving a maximum a posteriori (MAP) problem:

$$\hat{F_0} = \text{argmax} P(F_0|F,K) = \text{argmax} P(F|\hat{F_0},K)P(\hat{F_0}). \tag{3}$$

By assuming a Gaussian model and then taking its logarithmic energy function, the above MAP problem can be solved as the minimization of

$$E(\hat{F_0}|F,K) = \|\hat{F_0} \cdot K - F\|^2 + H(\hat{F_0}). \tag{4}$$

The regularization term $H(\hat{F_0})$ can be formulated using a variety of image priors. To simplify our analysis, we constrain $H(\hat{F_0})$ to be $\|C \cdot \hat{F_0}\|^2$, where C is a matrix. Then, minimizing $E(\hat{F_0}|F,K)$ gives us the well-known Wiener deconvolution [18]:

$$\hat{F_0} = \frac{F \cdot \bar{K}}{|K|^2 + |C|^2}, \tag{5}$$

where $\bar{K}$ is the complex conjugate of K, $|K|^2 = K \cdot \bar{K}$, and $|C|^2 = C \cdot \bar{C}$. Furthermore, the optimal $|C|^2$ is known to be the matrix of noise-to-signal ratios (NSR), $|\sigma/F_0|^2$.

We generally do not have access to the exact NSR matrix since F_0 is unknown. The traditional approach is to replace $|C|^2$ with a single scalar parameter λ or a simplified matrix like $\lambda \cdot (|G_x|^2 + |G_y|^2)$, where G_x and G_y are the Fourier transforms of the spatial derivative filters in the x-axis and y-axis, respectively. These simplifications cause deconvolution to not be optimal. More importantly, the parameter λ needs to be tuned, which is difficult as it is inherently scene dependent. Since we would like our aperture pattern evaluation/optimization to be automatic, we seek a deconvolution method that is free of parameter selection.

3.2. Optimizing Parameter C Using an Image Prior

Given a blur pattern K and a defocused image F, the focused image can be estimated as $\hat{F_0}$ by using Equation (5). Since noise ζ is a random matrix, we evaluate the quality of recovery using the expectation of the L_2 distance between $\hat{F_0}$ and the ground truth F_0 with respect to ζ:

$$R(K,F_0,C) = \mathop{\mathbb{E}}_{\zeta}[\|\hat{F_0} - F_0\|^2] = \mathop{\mathbb{E}}_{\zeta}\left\|\frac{\zeta \cdot \bar{K} - F_0 \cdot |C|^2}{|K|^2 + |C|^2}\right\|^2, \tag{6}$$

where $\mathbb{E}$ denotes expectation. When ζ is assumed to be Gaussian white noise $N(0,\sigma^2)$, we have

$$R(K,F_0,C) = \left\|\frac{\sigma \cdot \bar{K}}{|K|^2 + |C|^2}\right\|^2 + \left\|\frac{F_0 \cdot |C|^2}{|K|^2 + |C|^2}\right\|^2. \tag{7}$$

Since F_0 is sampled from the space of all images and has a certain distribution, we look for a C that minimizes the expectation of R with respect to F_0:

$$R(K,C) = \mathop{\mathbb{E}}_{F_0}[R(K,F_0,C)] = \int_{F_0} R(K,F_0,C)d\mu(F_0), \tag{8}$$

where $\mu(F_0)$ is the measure of the sample F_0 in the image space. According to the $1/f$ law of natural images [4][5][6], we know that the expectation of $|F_0|^2$,

$$A(\xi) = \int_{F_0} |F_0(\xi)|^2 d\mu(F_0), \tag{9}$$

exists (ξ is the frequency). Therefore, we can obtain

$$R(K,C) = \left\|\frac{\sigma \cdot \bar{K}}{|K|^2 + |C|^2}\right\|^2 + \left\|\frac{A^{1/2} \cdot |C|^2}{|K|^2 + |C|^2}\right\|^2. \tag{10}$$

For a given K, minimizing $R(C|K)$ gives us

$$|C|^2 = \sigma^2/A. \tag{11}$$

In practice, A can be estimated by simply averaging the power spectra of a lot of natural images.

3.3. Evaluating an Aperture Pattern

By substituting $|C|^2 = \sigma^2/A$ in Equation (10) and rearranging, we get the following metric that allows us to evaluate the quality of the aperture pattern K:

$$R(K) = \sum_{\xi} \frac{\sigma^2}{|K_\xi|^2 + \sigma^2/A_\xi}. \tag{12}$$

At each frequency ξ, $\frac{\sigma^2}{|K_\xi|^2 + \sigma^2/A_\xi}$ reflects the degree to which noise is amplified. The optimal pattern has the smallest $R(K)$.

Equation (12) highlights the fact that level of image noise σ is an important factor in evaluating an aperture pattern. It also suggests that, at different noise levels, the optimal aperture pattern can be different. It should be noted that this equation gives the expected performance of a pattern over the entire space of natural images, but might not be optimal for a given specific image. However, since the $1/f$ law is fairly robust, the optimized aperture patterns based on this criterion yield good deconvolution performances for a wide variety of real images.

4. Finding the Optimal Aperture Pattern

Even though our evaluation criterion (Equation (12)) is concise, finding the optimal aperture pattern remains a challenging problem. While the aperture pattern is evaluated in frequency domain, it must satisfy several physical constraints in spatial domain: (a) All its transmittance values must lie between 0 and 1; (b) the whole pattern should fit within the largest clear aperture of the lens; and (c) its spatial resolution must be low enough to avoid introducing strong diffraction effects. Deriving a closed-form optimal solution that satisfies all these constraints is difficult. We therefore resort to a numerical search approach. However, for a binary pattern of resolution $N \times N$, the number of possible solutions is $2^{N \times N}$, making exhaustive search impractical even for small values of N. In previous works that use other evaluation criteria [3] [13], randomized linear search has been used to find sub-optimal solutions.

We develop a genetic algorithm [7] to solve this optimization problem. We chose to use genetic algorithms as

Table 1. Genetic Algorithm for Aperture Pattern Optimization

1: Initialize: $g = 0$; randomly generate S binary sequences of length L.
2: For $g = 1 : G$
  a: *Selection*: For each sequence b, the corresponding blur function K is computed as $\Sigma P_{ij} * b_{i*N+j}$, and then evaluated by using Equation (12). Only the best M out of S sequences are selected.
  b: Repeat until the population (the number of sequences) increases from M to S.
    — *Crossover:* Duplicate two randomly chosen sequences from the M sequences of Step 2a, align them, and exchange each pair of corresponding bits with a probability of c_1, to obtain two new sequences.
    — *Mutation:* For each newly generated sequence, flip each bit with a probability c_2.
3: Evaluate all the remaining sequences using Equation (12) and output the best one.

* In our implementation, $L = 169$, $S = 4000$, $M = 400$, $c_1 = 0.2$, $c_2 = 0.05$ and $G = 80$.

they are known to rapidly find good solutions within complex binary search spaces. An aperture pattern k of size $N \times N$ can be expressed as $k = \sum_{i,j} p_{ij} \cdot b_{i \cdot N+j}$, where p_{ij} is a matrix, defined as

$$p_{ij}(x, y) = \begin{cases} 1, & \textit{for } [x, y] = [i, j] \\ 0, & \textit{otherwise} \end{cases},$$

b_{i*N+j} is 0 or 1, and $i, j \in [0, N-1]$.

Each aperture pattern is represented with a binary sequence b. In Fourier domain, we have $K = \Sigma P_{ij} * b_{i*N+j}$, where P_{ij} is the Fourier transform of p_{ij}. Note that p_{ij} should be zero-padded before computing the Fourier transform. The optimization can be sped up by pre-computing all P_{ij}. It is well-known in optics that an aperture of higher resolution will produce stronger diffraction effects. To this end, we set the spatial resolution $N \times N$ of our aperture function to be relatively low, i.e., N = 13.

The process of our genetic algorithm is described in Table 1. In our implementation, for a 13×13 pattern, a total of $S \times G = 320,000$ samples are evaluated, where S is the number of samples in each generation and G is the total number of generations. The optimization takes about 20 minutes on a 4GHz PC, and no significant improvement is observed with a larger G. We repeated the optimization ten times with different initial populations and found that it always converges to patterns with similar appearance.

As stated earlier, the optimal aperture pattern varies with the level of image noise. We performed our optimization using eight levels of noise; σ = 0.0001, 0.001, 0.002, 0.005, 0.008, 0.01, 0.02, and 0.03. The resulting apertures are shown in the bottom row of Figure 3. It is interesting to note that the optimized aperture patterns get simpler in structure with increase in noise.

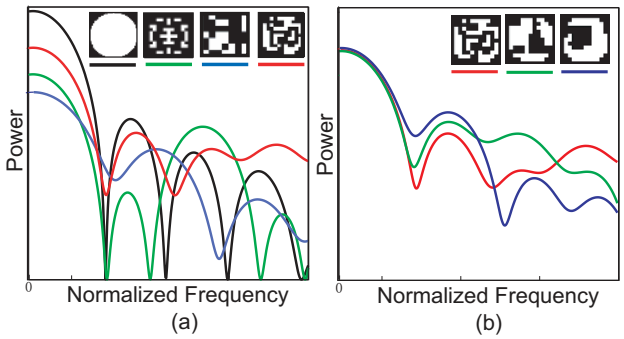

Figure 2. 1D slices of Fourier transforms of different patterns. (a) Circular pattern (black), Levin et al.'s pattern (green), Veeraraghavan et al.'s pattern (blue), and the optimized pattern for $\sigma = 0.001$ (red). (b) The optimized patterns for $\sigma = 0.001$ (red), $\sigma = 0.005$ (green), and $\sigma = 0.01$ (blue).

In Figure 2(a), we compare the power spectrum of one of our optimized apertures ($\sigma = 0.001$) with those of the circular pattern, Levin et al.'s pattern and Veeraraghavan et. al.'s pattern. Though these plots only show us 1D slices of 2D Fourier power spectra, they give us a strong intuition for how the various apertures would perform in the case defocus deblurring. Figure 2(a) shows that the circular pattern and Levin et al.'s pattern have many zero-crossings and greatly attenuate high frequencies. Again, it should be noted that Levin et. al.'s pattern is not designed for defocus deblurring. Veeraraghavan et al.'s pattern avoids zero-crossings, but it has lower response than our optimized aperture in both the low and high frequencies.

In Figure 2(b), we compare three of our optimized patterns ($\sigma = 0.001, 0.005, 0.01$). The optimized pattern for low noise has a larger response to high frequencies, while the one optimized for high noise has a larger response to low frequencies.

5. Deconvolution Algorithm

By substituting Equation (11), $|C|^2 = \sigma^2/A$, into Equation (13), we obtain the following variant of Wiener deconvolution :

$$\hat{F_0} = \frac{F \cdot \bar{K}}{|K|^2 + \sigma^2/A}. \tag{13}$$

Note that this deconvolution algorithm is optimal in the sense of minimizing the expected L_2 distance between the deblurred image and the ground truth. Its results, though potentially less visually appealing than methods using sparse priors, can be expected to be more faithful to the ground truth in this sense. More importantly, the matrix A as defined in Equation (9) can be estimated by simply averaging the power spectra of several natural images, and the noise level σ can be approximated from the model of the camera and its ISO (or gain) setting. Consequently, in contrast to most other deconvolution methods, this deblurring algorithm is free of parameter tuning. For these reasons, we

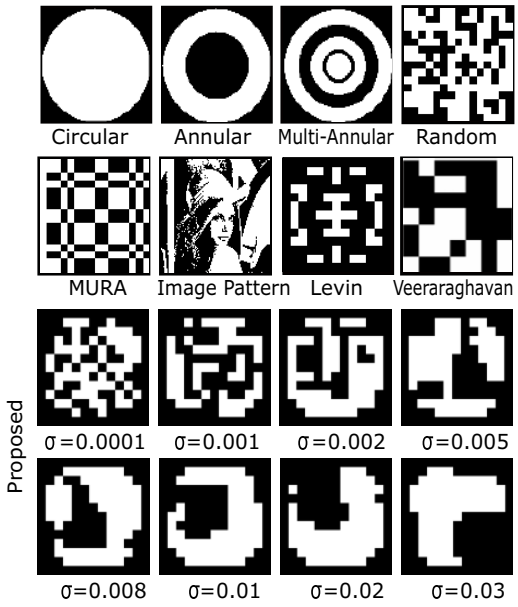

Figure 3. All the aperture patterns we used in our simulations. Top two rows: Eight patterns, including circular, annular, multi-annular, random, MURA, image pattern, Levin et al.'s pattern [13], and Veeraraghavan et al.'s pattern [3]. Bottom two rows: Eight of our patterns optimized for noise levels from $\sigma = 0.0001$ to 0.03.

have used it in all of our comparisons and experiments. It must be noted that similar algorithms have been advocated in the past (see [19] for example).

6. Performance Comparison of Apertures

Before conducting real experiments, we first performed extensive simulations to verify our aperture evaluation criterion and optimization algorithm. For this, we used the 16 aperture patterns shown in Figure 3. The top 8 patterns include simple ones (circular, annular, and multi-annular) and more complex ones proposed by other researchers [9], [13], [3]. In addition, we have tested an "image pattern," which is a binarized version of the well-known Lena image, and a random binary pattern. The bottom 8 patterns were produced by our optimization algorithm for different levels of image noise. The performances of these 16 apertures were evaluated via simulation over a set of 10 natural images at eight levels of image noise.

For each aperture pattern k and each level of image noise σ, we simulated the defocus process using Equation (1), applied defocus deblurring using Equation (13), and got an estimate $\hat{f}_0$ of the focused image f_0. Using each deblurred image, the quality of the aperture pattern was measured as $\sqrt{\|f_0 - \hat{f}_0\|^2}$. To make this measurement more reliable, we repeated the simulation on 10 natural images and took the average. These results are listed in Table 2 for the 16 aperture patterns and 8 levels of image noise. Our optimized patterns perform best across all levels of noise, and the improvement is more significant when the noise level is low. On the other hand, the circular (conventional) aperture is close to optimal when the noise level is very high. While there are different optimal apertures for different levels of image noise, we may want a single aperture to use in a variety of imaging conditions. In this case, we could pick the optimized pattern for $\sigma = 0.001$ as it performed well over a wide range of noise levels (from $\sigma = 0.0001$ to 0.01).

It is interesting to note that the image pattern (Lena) also produces deblurring results of fairly high quality. We believe this is because the power spectrum of the image pattern follows the $1/f$ law–it successfully avoids zero-crossings and, at the same time, has a heavy tail covering the high frequencies. Unfortunately, the image pattern consists of a lot of small features, which introduce strong diffraction effects. We believe that it is for this reason that the image pattern did not achieve as high quality results in our experiments as predicted by our simulations.

7. Experiments with Real Apertures

As shown in Figure 4(a), we printed our optimized aperture patterns as well as several other patterns as a single high resolution (1 micron) photomask sheet. To experiment with a specific aperture pattern, we cut it out of the photomask sheet and inserted it into a Canon EF $50mm\ f/1.8$ lens[1]. In Figure 4(b), we show 4 lenses with different apertures (image pattern, Levin et al.'s pattern, Veeraraghavan et al's pattern, and one of our optimized patterns) inserted in them, and one unmodified (circular aperture) lens. Images of real scenes were captured by attaching these lenses to a Canon EOS 20D camera. As previously mentioned, we choose the pattern which is optimized for $\sigma = 0.001$, as it performs well over a wide range of noise levels in the simulation.

To calibrate the true PSF of each of the 5 apertures, the camera focus was set to $1.0m$; a planar array of point light sources was moved from $1.0m$ to $2.0m$ with $10cm$ increments; and an image was captured for each position. Each defocused image of a point source was deconvolved using a registered focused image of the source. This gave us PSF estimates for each depth (source plane position) and several locations in the image[2]. In Figure 4(c-g), two calibrated PSFs (for depths of $120cm$ and $150cm$) are shown for each pattern.

7.1. Comparison Results using Test Scenes

In our first experiment, we placed a CZP resolution chart at a distance of $150cm$ from the lens, and captured images using the five different apertures. To be fair, the same exposure time was used for all the acquisitions. The five captured images and their corresponding deblurred results are shown

[1] We chose this lens for its high quality and because we were able to disassemble it to insert aperture patterns with relative ease.

[2] We measured the PSF at different image locations to account for the fact that virtually any lens (even with a circular aperture) produces a spatially varying PSF.

Table 2. Performance comparison of 16 aperture patterns for eight noise levels.

Patterns	**Image Noise Level σ**							
	0.0001	0.0005	0.001	0.002	0.005	0.008	0.01	0.02
Circular	0.0234	0.0375	0.0439	0.0503	0.0587	0.0631	0.0652	0.0717
Annular	0.0194	0.0334	0.0405	0.0478	0.0573	0.0622	0.0645	0.0716
Multi-Annular	0.0141	0.0274	0.0346	0.0426	0.0537	0.0598	0.0627	0.0719
Random	0.0157	0.0294	0.0368	0.0448	0.0558	0.0616	0.0645	0.0731
MURA	0.0153	0.0279	0.0345	0.0419	0.0531	0.0594	0.0624	0.0719
Image pattern	0.0128	0.0252	0.0324	0.0403	0.0513	0.057	0.0597	0.0681
Levin	0.0181	0.0316	0.0394	0.0486	0.0619	0.0686	0.0716	0.0798
Veeraraghavan	0.0164	0.0282	0.0346	0.0419	0.0527	0.0586	0.0614	0.0703
Optimized Patterns for:								
$\sigma = 0.0001$	**0.0118**	**0.0235**	0.0313	0.0407	0.0544	0.0613	0.0644	0.0732
$\sigma = 0.001$	0.0123	0.024	**0.0309**	**0.039**	0.0513	0.0581	0.0614	0.0713
$\sigma = 0.002$	0.0135	0.0261	0.0327	0.0398	**0.0501**	**0.0561**	0.059	0.0686
$\sigma = 0.005$	0.0138	0.0269	0.034	0.0415	0.0513	**0.0561**	**0.0585**	0.0663
$\sigma = 0.008$	0.014	0.0276	0.035	0.0425	0.052	0.0566	0.0588	**0.0659**
$\sigma = 0.01$	0.0144	0.028	0.0353	0.043	0.0527	0.0572	0.0593	**0.0659**
$\sigma = 0.02$	0.0151	0.029	0.0366	0.0447	0.0548	0.0593	0.0612	0.0671
$\sigma = 0.03$	0.0157	0.0301	0.0377	0.0454	0.055	0.0594	0.0614	0.0674

** The best performer for each noise level is shown in bold.*

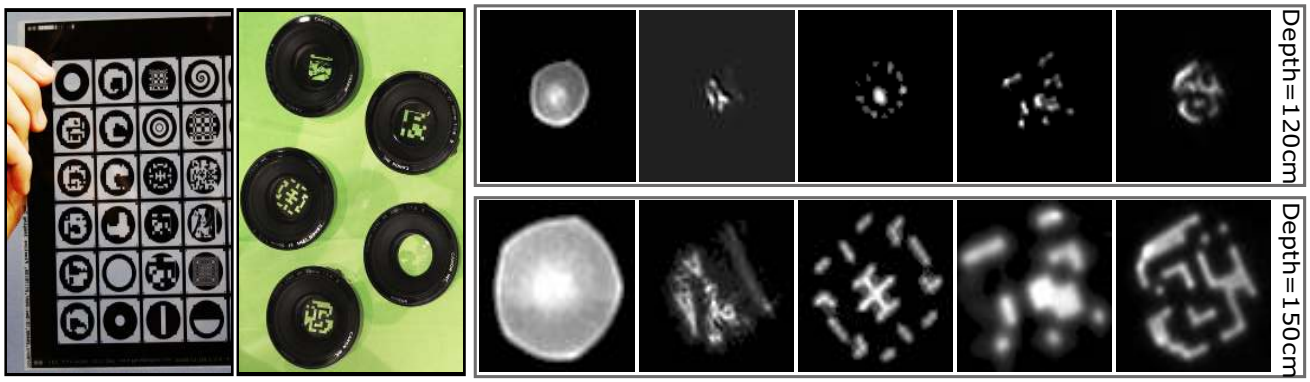

Figure 4. (a) Photomask sheet with many different aperture patterns. (b) One unmodified lens and four lenses with patterns inserted. (c-g) Top row shows calibrated PSFs for a depth of $120cm$ from the lens, and bottom row shows calibrated PSFs for a depth of $150cm$. These PSFs correspond to (c) circular pattern, (d) image pattern, (e) Levin et al.'s pattern, (f) Veeraraghavan et al.'s pattern, and (g) one of our optimized patterns.

in Figures 1 and 5. Notice that the captured images have different brightness levels as the apertures obstruct different amounts of light. The resulting brightness drop (compared to the circular aperture) for the image pattern, Levin et al.'s pattern, Veeraraghavan et al.'s pattern, and our optimized pattern are 48%, 52%, 65%, and 43%, respectively.

Note that our optimized pattern gives the sharpest deblurred image with least artifacts and image noise (see Figures 1 and 5). We performed a quantitative analysis to compare the performances of the five apertures. We carefully aligned all the deblurred images to the focused image with sub-pixel accuracy, and computed their residual errors. The residual errors were then analyzed in frequency domain. In Figure 5(d), we plot the cumulative energy of the residual error from low to high frequency. The image pattern, Levin et al.'s pattern, and especially Veeraraghavan et al.'s pattern, show large improvements over the circular aperture. Our optimized aperture is seen to produce the lowest residual error with about 30% improvement over Veeraraghavan et al.'s pattern (which performs the best among the rest).

7.2. Deblurring Results for Complex Scenes

We have used the lens with our optimized aperture pattern to capture several real scenes with severely defocused regions (see Figure 6). Deblurring of a region requires prior knowledge of its depth. In all our examples, we interactively selected the depth that produced the most appealing deblurring result. This is made possible by the fact that our deblurring algorithm described in Section 3.2 is very fast and requires no parameter selection. For a 1024×768 image, our Matlab implementation of the algorithm takes only about 100msec to run. In contrast, state-of-the-art deblurring algorithms, such as ones that use sparse priors, are much slower and require the selection of parameters.

Figure 6(a) shows a captured image (left) for which the camera was focused on the foreground object, making the background poster severely defocused. To deblur the back-

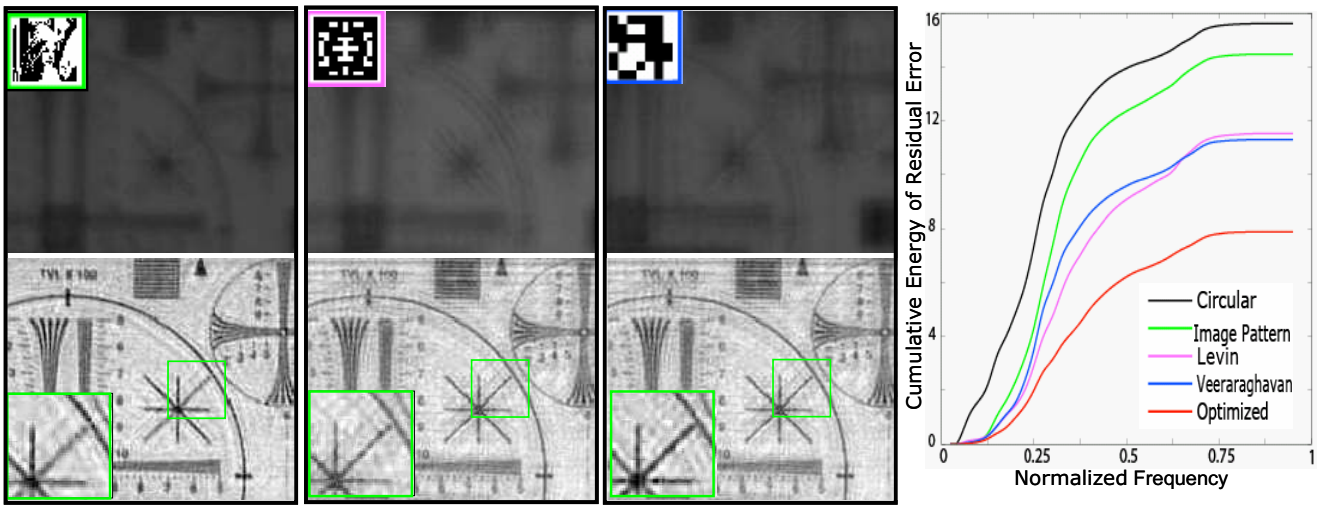

(a) Image Pattern (b) Levin (c) Veeraraghavan (d) Cumulative Residual Energy

Figure 5. (a-c) The top row shows captured (defocused) images and the bottom row shows the deblurred images, for three different apertures. The focused image (ground truth) and the results using the circular aperture and our optimized aperture are shown in Figure 1. (d) For each aperture, the cumulative energy of the residual error between the ground truth and deblurred images is plotted as a function of frequency.

ground, we first segmented out the foreground region, filled the resulting hole using inpainting, and then applied deblurring using 40 different depths. The best deblurred result is chosen and merged with the foreground. Figure 6(b) shows a traffic scene where all the objects are out of focus. In this case, the final result was obtained using four depth layers. Although some ringing artifacts can be seen in our deblurred images, significant details are recovered in all cases. It may be noted that the degree of defocus in our experiments is much greater than in the experiments done in previous works [13][3]. For example, the recovered telephone number and taxi number in Figure 6(b) are virtually invisible in the captured image.

8. Discussion

In this work, we presented a comprehensive criterion for evaluating aperture patterns for the purpose of defocus deblurring. This criterion explicitly accounts for the effects of image noise as well as the statistics of natural images. To make the aperture pattern optimization tractable, we have assumed a Gaussian white noise model. This noise model may not be accurate in some imaging systems, and could result in sub-optimal solutions. Enabling the use of more elaborate noise models and making use of an even stronger image prior in the aperture optimization are interesting directions we plan to pursue in future work.

Diffraction is another important issue that requires further investigation. Our work, as well as previous works on coded apertures, have avoided having to deal with diffraction by simply using low-resolution aperture patterns. By explicitly modeling diffraction effects, we may be able to find even better aperture patterns for defocus deblurring.

References

[1] W. Welford, "Use of annular apertures to increase focal depth," *Journal of the Optical Society of America A*, no. 8, pp. 749–753, 1960. 1, 2

[2] M. Mino and Y. Okano, "Improvement in the OTF of a defocused optical system through the use of shaded apertures," *Applied Optics*, no. 10, pp. 2219–2225, 1971. 1, 2

[3] A. Veeraraghavan, R. Raskar, A. Agrawal, A. Mohan, and J. Tumblin, "Dappled photography: mask enhanced cameras for heterodyned light fields and coded aperture refocusing," *ACM Transactions on Graphics*, 2007. 1, 2, 3, 5, 7

[4] D. Mumford and B. Gidas, "Stochastic models for generic images," *Quarterly of Applied Mathematics*, no. 1, pp. 85–111, 2001. 1, 3

[5] A. Srivastava, A. Lee, E. Simoncelli, and S. Zhu, "On Advances in Statistical Modeling of Natural Images," *Journal of Mathematical Imaging and Vision*, pp. 17–33, 2003. 1, 3

[6] Y. Weiss and W. Freeman, "What makes a good model of natural images?" *CVPR*, pp. 1–8, 2007. 1, 3

[7] M. Srinivas and L. Patnaik, "Genetic algorithms: a survey," *Computer*, no. 6, pp. 17–26, 1994. 1, 3

[8] E. Caroli, J. Stephen, G. Cocco, L. Natalucci, and A. Spizzichino, "Coded aperture imaging in X- and Gamma-ray astronomy," *Space Science Reviews*, pp. 349–403, 1987. 2

[9] S. Gottesman and E. Fenimore, "New family of binary arrays for coded aperture imaging," *Applied Optics*, no. 20, pp. 4344–4352, 1989. 2, 5

[10] E. Dowski and W. Cathey, "Extended depth of field through wave-front coding," *Journal of the Optical Society of America A*, no. 11, pp. 1859–1866, 1995. 2

(a) Indoor Scene

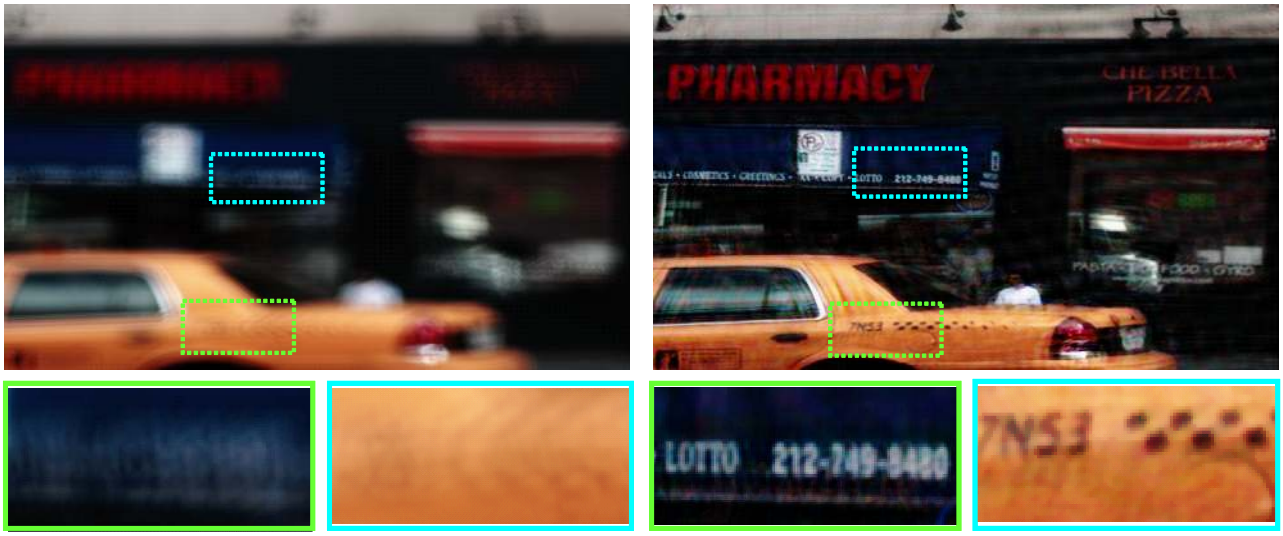

(b) Traffic Scene

Figure 6. Deblurring results for two complex scenes. Left: Captured images with close-ups (green and blue boxes) of regions that are severely defocused. Right: The corresponding deblurring results.

[11] N. George and W. Chi, "Extended depth of field using a logarithmic asphere," *Journal of Optics A: Pur and Applied Optics*, 2003. 2

[12] H. Nagahara, S. Kuthirummal, C. Zhou, and S. Nayar, "Flexible Depth of Field Photography," *ECCV*, 2008. 2

[13] A. Levin, R. Fergus, F. Durand, and W. Freeman, "Image and depth from a conventional camera with a coded aperture," *ACM Transactions on Graphics*, no. 3, 2007. 2, 3, 5, 7

[14] A. Zomet and S. Nayar, "Lensless imaging with a controllable aperture," *CVPR*, pp. 339–346, 2006. 2

[15] M. Aggarwal and N. Ahuja, "Split Aperture Imaging for High Dynamic Range," *International Journal of Computer Vision*, vol. 58, no. 1, pp. 7–17, 2004. 2

[16] P. Green, W. Sun, W. Matusik, and F. Durand, "Multi-aperture photography," *ACM Transactions on Graphics*, vol. 26, no. 3, 2007. 2

[17] C. Liang, T. Lin, B. Wong, C. Liu, and H. Chen, "Programmable aperture photography: Multiplexed light field acquisition," *ACM Transactions on Graphics*, vol. 27, 2008. 2

[18] H. Andrews and B. Hunt, "Digital image restoration," *Prentice-Hall Signal Processing Series, Englewood Cliffs: Prentice-Hall*, 1977. 3

[19] S. Reeves, "Image deblurring - wiener filter," Matlab Central Blog, http://blogs.mathworks.com/steve/2007/11/02/image-deblurring-wiener-filter/, November 2007. 5