Design and Estimation of Coded Exposure Point Spread Functions

Scott McCloskey, *Member*, *IEEE*,
Yuanyuan Ding, *Member*, *IEEE*, and
Jingyi Yu, *Member*, *IEEE*

Abstract—We address the problem of motion deblurring using coded exposure. This approach allows for accurate estimation of a sharp latent image via well-posed deconvolution and avoids lost image content that cannot be recovered from images acquired with a traditional shutter. Previous work in this area has used either manual user input or alpha matting approaches to estimate the coded exposure Point Spread Function (PSF) from the captured image. In order to automate deblurring and to avoid the limitations of matting approaches, we propose a Fourier-domain statistical approach to coded exposure PSF estimation that allows us to estimate the latent image in cases of constant velocity, constant acceleration, and harmonic motion. We further demonstrate that previously used criteria to choose a coded exposure PSF do not produce one with optimal reconstruction error, and that an additional 30 percent reduction in Root Mean Squared Error (RMSE) of the latent image estimate can be achieved by incorporating natural image statistics.

Index Terms—Coded exposure, motion deblurring, blur estimation, computational photography.

1 Introduction

Motion blur adversely impacts aesthetics and reduces the utility of images for automated analysis. Many algorithms mitigate the effects of motion blur in traditionally acquired images, but fundamental limits to the quality of the reconstruction are imposed by the use of a traditional shutter. Techniques that hallucinate or otherwise use information (priors, etc.) from outside of the image may produce visually pleasing images, but have little value for forensic applications, e.g., biometrics, where knowledge about the true scene texture is required.

Introducing coded exposure photography, Raskar et al. [1] point out that the lost frequencies introduced by traditional imaging of moving objects can be avoided by fluttering the shutter in order to capture an image with a motion blur PSF that can be removed by well-posed deconvolution. This improves the performance of forensic image exploitation, but requires new techniques for automated processing of coded exposure images.

Using natural image statistics, we introduce a blur estimation technique in Section 6 for automated analysis and deblurring of coded exposure images. Our approach uses the telltale inflection points of coded exposure Modulation Transfer Functions (MTFs) which appear in the amplitude spectra of captured images. Using this approach, we avoid specialized assumptions about the spatial content of images and demonstrate estimation of a wider range of motions. We also use natural image statistics in Section 5 to select coded exposure PSFs providing lower reconstruction error than the existing method.

In regions of uniform blur (such as a translating planar surface seen from a stationary camera), we model the blurred image I by convolution as

$$I = J \otimes B + \eta, \tag{1}$$

where J is the latent sharp image, B is the blur PSF, and η is imaging noise. As with other work using this model, we are limited to blurred images with spatially invariant PSFs unless motion segmentation [2] is performed as a preprocess. In addition, we are limited to motions with linear trajectories or to those such as harmonic motion that can be linearized. As such, our methods are intended to capture moving objects from stationary cameras, not for camera motion producing 2D PSFs.

2 Related Work

Our work is an extension of the coded exposure technique proposed in [1]. Though the original work depended on manual blur estimation by a user, Agrawal and Xu [3] address coded exposure blur estimation for constant velocity based on alpha matting, and choose exposure sequences for both invertibility and ease of blur estimation. As compared to [1] and [3], we improve the selection of the coded exposure PSF, and propose a Fourier-domain blur estimation technique which avoids the alpha matting assumption that there exists a high-contrast edge between moving and stationary objects. Coded exposure is an example of computational photography, where image capture and postprocessing are codesigned to improve image quality. Zhou and Nayar provide an excellent survey in [4].

Since most images are captured using traditional cameras, numerous blind deconvolution methods [5] help mitigate the effects of blur in traditional images. Spatially varying blur has attracted attention recently, with methods by Whyte et al. [6] and Gupta et al. [7]. Other methods assume constant velocity motion of planar surfaces along a straight line, in which case the PSF is a spatially invariant 1D box filter. Levin [2] examines the effects of such PSFs on image statistics to segment and deconvolve blurred regions. Where box PSFs lose certain spatial frequencies, outside information in the form of gradient or edge priors [8], [9], [10], [11], [12] help derive visually pleasing images.

Whether for deblurring, motion recovery, or depth recovery [13], estimating a blur PSF is critical to many methods. PSF estimation from a single image being ill posed, methods make it tractable in different ways. Some operate on multiple images taken with different settings, e.g., [13]. Others make assumptions about the scene, relying on priors such as color or strong edges [14], [15], [16]. Still others constrain the solution space to constant velocity, e.g., Cepstrum methods [17], [18], characterize the motion by the number and position of zeros in an image's power spectrum. These methods cannot be applied to coded exposure images that are acquired specifically to avoid such zeros. As in [3], Dai and Wu [19] treat motion blur as an alpha matte for estimating the PSF and implicitly assume the existence of high-contrast edges in the latent sharp image. Moreover, since matting only estimates the blur *extent*, it cannot distinguish motion direction or the infinite number of velocity/acceleration combinations that might produce that extent, which our method provides.

3 Conventions and Notation

Since a fluttering shutter is closed for part of the capture time, we distinguish between the *exposure time* (how long the shutter is open during capture) and the *capture time* (from the first shutter open to the last shutter close, including closed shutter periods). Though the exposure and capture times of a traditional shutter are equal, the flutter shutter's capture time is often twice the exposure time.

- *S. McCloskey is with Honeywell ACS Labs, 1985 Douglas Drive North #112A, Golden Valley, MN 55422. E-mail: scott.mccloskey@honeywell.com.*
- *Y. Ding is with Epson Research and Development, Inc., 214 Devcon Drive, San Jose, CA 95112. E-mail: yding@erd.epson.com.*
- *J. Yu is with the Department of CIS, University of Delaware, 101 Smith Hall, Newark, DE 19716. E-mail: yu@eecis.udel.edu.*

Manuscript received 3 Oct. 2011; revised 27 Jan. 2012; accepted 30 Apr. 2012; published online 8 May 2012.
Recommended for acceptance by C.-K. Tang.
For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-2011-10-0703.
Digital Object Identifier no. 10.1109/TPAMI.2012.108.

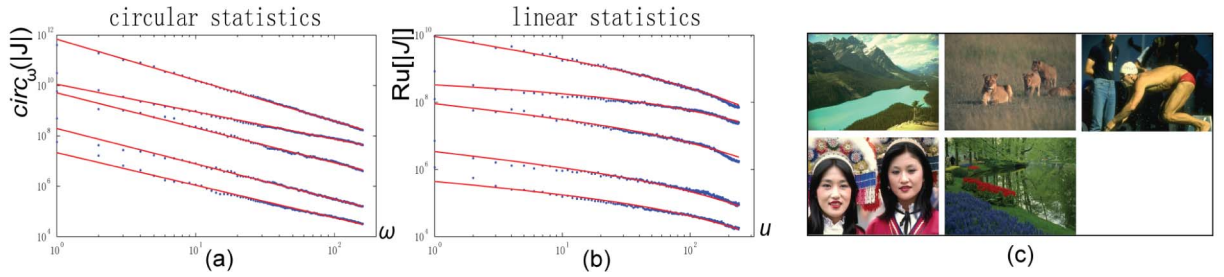

Fig. 1. Power spectrum statistics of five natural images (c). (a) The circular power spectrum versus the spatial frequency magnitude ω in a log-log scale. The red lines show the fits of the $1/\omega$-exponent model. (b) The linear statistics along v versus u in a log-log scale. The red curves show our estimated linear statistics from the circular statistics. For clarity, traces in both plots are shifted -1, -2, -3, and -4 log units.

When comparing traditional and flutter shutter images, we present images with equal exposure times.

Latent image estimates presented here have been deblurred using the reference code from [1]. We refer to the coded PSF given in [1] as Raskar's sequence.

By common convention, we denote image domain quantities (e.g., the blur PSF) as characters such as B and Fourier domain quantities (e.g., the blur Optical Transfer Function) as the corresponding character with a hat, such as $\hat{B}$. Of particular interest is the MTF, denoted $|\hat{B}|$, with optional parameter $|\hat{B}(k)|$, denoting a particular spatial frequency k.

4 Natural Image Statistics

Van der Schaaf and van Hateren [20] have shown that the power spectrum for a natural image J (without motion blur) falls off with absolute frequency. Thus, if we parameterize $|\hat{J}|$ in polar coordinates (ω, ϕ) (where ω is the frequency magnitude and ϕ is the angle), then average $|\hat{J}|$ over ϕ, the resulting circularly averaged spectrum $circ_\omega(|\hat{J}|) \approx \frac{C}{\omega^m}$, where m and C are constants. Assuming every frequency to be independent and identically distributed, the expected value of $|\hat{J}(u,v)|$ is

$$E[|\hat{J}(u,v)|] = \frac{C}{(u^2+v^2)^{m/2}}. \tag{2}$$

Fig. 1a shows circular spectra of (and fits to) five images from the Berkeley Segmentation Data set (BSD) [21].

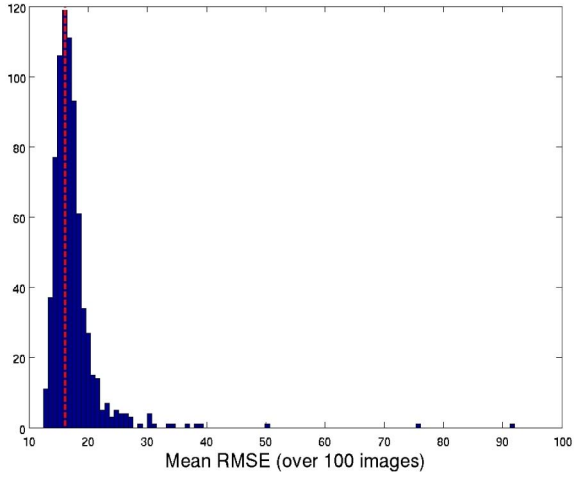

Fig. 2. RMSEs of latent image estimates using different PSFs. The histogram shows the average RMSE over 100 images using the traditional shutter ($\mathrm{RMSE} = 75$) and 749 equal exposure coded PSFs. The PSF satisfying the max-min criterion has $\mathrm{RMSE} = 16$ (red dashed line), about 30 percent worse than the best.

Because we consider 1D blur due to linear motion, we replace circular power spectrum statistics with linear statistics. Specifically, we *project* the 2D amplitude spectrum onto a line l in the motion direction by rotating the spectrum so that l is aligned with the u-axis and integrating over v. In the discrete case, we compute

$$R_u[|\hat{J}|] = \frac{1}{V}\sum_{v=0}^{V} |\hat{J}(u,v)|, \tag{3}$$

where V is the v-dimension resolution. $R_u[|\hat{J}|]$ represents the horizontal power spectrum statistics and can be approximated using (2) as

$$R_u[|\hat{J}|] \approx E\left[\frac{1}{V}\sum_{v=0}^{V} |\hat{J}(u,v)|\right] = \frac{1}{V}\sum_{v=0}^{V} \frac{C}{(u^2+v^2)^{m/2}}. \tag{4}$$

Fig. 1 shows that this linear model (b) has similar accuracy to the circular model (a).

5 Determining the Shutter Sequence

In this section, we demonstrate that natural image statistics can be used to select a PSF which provides better latent image estimates than the criteria used to find Raskar's sequence. The flutter shutter objective is to acquire images of moving objects that, while blurred, have a motion blur PSF which is invertible by deconvolution. While it is known that lost spatial frequencies are problematic, we must choose between those PSFs that avoid lost frequencies.

In [1], random PSFs are sampled, and one is chosen to have its Fourier transform (the MTF) satisfy $argmax_B \min_k \|\hat{B}(k)\|$.[1] We demonstrate the inadequacy of this max-min criterion using 750 equal exposure PSFs: a traditional shutter PSF with 42 open shutter chops and 749 random permutations thereof. For each of the 750 PSFs, we synthetically blurred the 100 test images from the BSD, added noise (Gaussian with $\sigma = 1$ relative to intensities in [0, 255]), deblurred, and computed the RMSE versus the original image. Fig. 2 shows the histogram of mean RMSE values for these PSFs, and illustrates two important points. First, *even a random coded exposure PSF will give much lower RMSE than a traditional shutter*. Second, it shows that the max-min criterion is insufficient, as the PSF satisfying it produces an RMSE 30 percent worse than the optimal, comparable to random selection.

5.1 Natural Image Criteria

The failure of the max-min criterion to produce the best PSF can be understood by recalling that natural images have rapidly diminishing power in higher spatial frequencies. The max-min criterion selects for a contrast-preserving PSF *without differentiating based on the absolute frequency corresponding to the minimum*, despite

1. Per the code available here: www.umiacs.umd.edu/~aagrawal/MotionBlur/SearchBestSeq.zip.

TABLE 1
Correlations between Computed Metrics and Observed RMSE of Deblurred Images

Metric	Correlation with RMSE	Probability	Weight
m_1	-0.0288	0.0125	0.4076
m_2	-0.2489	$\ll .05$	-255.2645
m_3	-0.3527	$\ll .05$	-1.2188e+03
m_4	0.3527	$\ll .05$	-2.4768e+03
m_5	0.0586	$\ll .05$	-778.1234
m_6	0.2398	$\ll .05$	1.1007
m_7	0.1438	$\ll .05$	-5.4183e+05

The center column shows the probability that the correlation results from random chance; values less than 0.05 are generally considered to indicate a significant effect. The right column shows the weight used in (5).

natural images having highly uneven power distribution. Since natural images have more power in low frequencies, we improve RMSE performance by finding PSFs which preserve higher contrast in low frequency bands.

In addition to the MTF minimum, we computed a number of other metrics on the 750 PSFs. The complete list of metrics collected is

- **Number of open shutter periods** m_1,
- **MTF Minimum** $m_2 = min_k|\hat{B}(k)|$,
- **MTF Mean** $m_3 = mean_k(|\hat{B}(k)|)$,
- **MTF Variance** $m_4 = var_k(|\hat{B}(k)|)$,
- **Prior-weighted mean** $m_5 = mean_k \frac{|\hat{B}(k)|}{k+1}$,
- **Number of lost frequencies** $m_6 = \|\{k \; where \; |\hat{B}(k)| < 10^{-10}\}\|$,
- **Weighted lost frequencies** $m_7 = \sum_k \frac{\chi(k)}{k+1}$, where $\chi(k) = 1 \Longleftrightarrow |\hat{B}(k)| < 10^{-10}$.

The correlations between these metrics and the PSF's mean RMSE (over the 100 images) are shown in Table 1. In order to use these metrics to find the optimal PSF, we learn a function mapping a vector of the metrics to average RMSE. With the metrics and RMSE computed for the 750 PSFs, we use multiple regression (Matlab's \ operator) to find weights w such that the quantity

$$\sum_{i=0}^{7} w_i m_i, \tag{5}$$

best fits the mean RMSE, with $m_0 = 1$ included to provide an offset term.[2] Our weights are given in Table 1, and $w_0 = 1{,}446$. The resulting weighted combination of metrics improves on the base metrics, providing a slightly higher correlation of 0.4408 to RMSE. While this appears to be a modest improvement, we will demonstrate in the next section that even such a weakly predictive projection can be used to select a PSF that provides near-optimal RMSE.

5.2 PSF Design Algorithm

Given the experiments presented in the previous section, we propose to use a combination of natural image-weighted metrics and synthetic blurring/deblurring to choose a PSF that provides low reconstruction RMSE. The high-level steps of our algorithm are as follows:

1. First, we find the actual RMSE for a small number of PSFs over a representative set of images.
2. Using these actual RMSEs and computed metrics on the PSFs, we determine weights by regression.
3. We then use that weighting and (5) to project the RMSEs from the metrics of a large number of other randomly sampled PSFs, of which we retain those candidates with the lowest projected RMSEs.
4. Finally, we find the actual RMSEs over the images for the PSFs in the small set of candidates, and choose the one with the lowest actual RMSE.

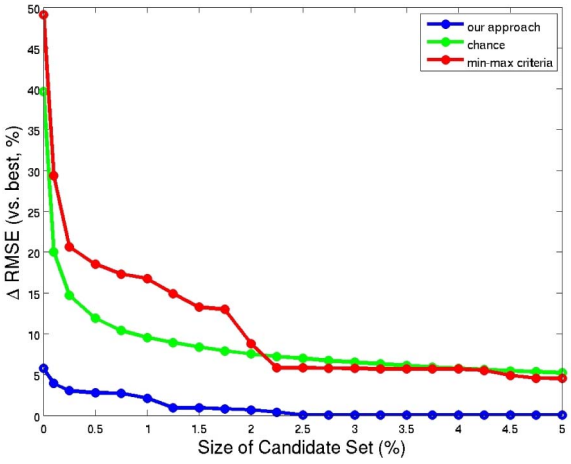

Fig. 3. With natural image statistics, we select a PSF providing near-optimal reconstruction while finding the actual RMSE on a small subset of the search space. As compared to chance and max-min sampling, the proposed algorithm produces significantly better results. Note: The leftmost points lie at $\frac{1}{50}$, not 0.

In order to evaluate this algorithm's effectiveness, we selected an additional 7,500 equal-exposure PSFs. For each, we find both the metrics listed above and the ground truth RMSE over the 100 BSD images.[3]

From this data, we simulate 1,000 runs of our algorithm. For each, we randomly select 5,000 of the 7,500 PSFs as the search space and use our method to find the best one. We compare the actual RMSE of that PSF to the one with the ground truth optimal RMSE in the space. Ideally, the difference would be 0, indicating that we found the best PSF of those in the search space. Fig. 3 shows a plot of this difference as a function of the size of the candidate set. Over the 1,000 trials, the average RMSE of the best PSF was 12.00 and the average RMSE of the PSF with the best weighted combination of metrics was 12.74 (difference of 0.74 or 6.2 percent), indicating that, while the weighted metrics outperform the max-min criterion, the projection does *not* generally produce the optimum PSF. When the 1 percent of the PSFs with the best projected RMSE are evaluated to find their actual RMSE, the difference narrows to 0.26 (2.2 percent), and the difference is further reduced as the candidate set is enlarged (the blue curve). For comparison, we show that if the candidate set is comprised of randomly selected PSFs instead of using our metric weighting, the comparable algorithm would produce the green curve. Likewise, if the candidate set were formed using the max-min criterion, the algorithm would produce the red curve. Our algorithm produces an actual RMSE that is closer to that of the optimal PSF, significantly improving on [1].

6 Coded Exposure Blur Estimation

In this section, we motivate and develop a method to estimate the PSF from a single coded exposure image. To handle both accelerated and harmonic motion, we must estimate more than just the width of a motion-blurred edge since, for example, there are infinitely many pairs of starting velocities and accelerations

2. In our experiments, we found that including second-order products occasionally provided a slight benefit; higher-order products of the metrics do not improve the quality of the fit.

3. Note that this takes approximately 1 CPU week, limiting the number of PSFs used in our experiment. At this rate, exhaustively searching all potential PSFs would take more than 10^9 years.

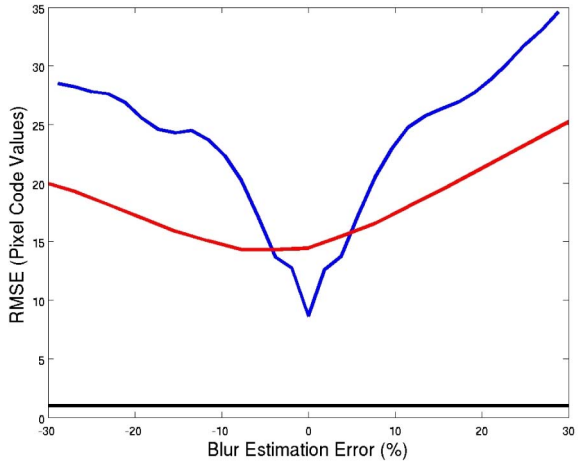

Fig. 4. Latent image estimation RMSE as a function of blur estimation error. (Blue) RMSE for deblurred coded exposure images. (Red) RMSE for deblurred traditional shutter images.

that will produce a given width. Because of this, and to avoid assuming that the scene contains a high-contrast edge between the moving object and background, we do not base our approach on alpha matting. Instead, we assume that the image contains only the motion-blurred object[4] with texture obeying a natural image prior. Relative to such a prior, we note that the local minima and maxima of a coded exposure MTF introduce features in the image's amplitude spectrum that can be used to estimate motion parameters. Before presenting our method, however, we first quantify the sensitivity of coded exposure deblurring to errors in PSF estimation.

6.1 Sensitivity to Blur Estimation Error

Fig. 4 shows latent image estimation RMSE as a function of the accuracy of PSF length estimates for both traditional and coded exposure images. We synthesized traditional and coded exposure motion blur in 10 randomly selected images from the BSD, and added Gaussian noise η with $\sigma = 1$. After deblurring, we take the mean RMSE of the 10 images. The plot shows that coded exposure deblurring, whose RMSE rises sharply around the minimum, is more sensitive than traditional imaging to errors in blur estimation. The rapid RMSE increase for coded exposure deblurring makes image quality a good proxy for PSF accuracy, which is important because establishing ground truth is difficult for real moving objects. When blur estimation error is more than ± 5 percent, latent images estimated from traditional shutters are quantitatively better than those using coded exposure.

We stress that the quantitative results shown in Fig. 4 should only serve as a rough guide to the differences between traditional and coded exposure imaging. At points where the red and blue curves intersect, differences in the MTFs will result in significant qualitative differences. Fig. 5 shows estimated latent images having nearly identical RMSEs, but notable qualitative differences: The coded exposure images are sharp but noisy, whereas the traditional shutter images have soft focus but low noise.

6.2 PSF Modeling

The unit-area PSF $B(x)$ describes how long each pixel x is exposed to a moving scene point. It is a function of both the binary shutter sequence $S(t)$ and the object motion α, which we measure in units of pixels, e.g., velocity as pixel/sec. A pixel x is exposed to a moving point with duration $w(x)$ inversely proportional to velocity $\nu(x)$ as

$$w(x) = \frac{1}{\nu(x)}. \tag{6}$$

It is natural to describe the velocity and displacement in terms of time t, so we rewrite $w(x) = \frac{1}{\nu(t(x))}$, where $t(x)$ is the inverse of the displacement function $x(t)$. We assume that $x(t)$ is monotonic throughout the shutter sequence, i.e., there is no back and forth motion.

We combine the shutter sequence and the exposure $w(x)$ to compute the unnormalized PSF $B_0(x)$ as

$$B_0(x) = S(t(x))w(t(x)) = \frac{S(t(x))}{\nu(t(x))}. \tag{7}$$

$B_0(x)$ is then normalized to the unit-area PSF $B(x)$. Equation (7) indicates that the PSF can be viewed as an envelope of $w(x)$ sampled by the shutter pattern $S(t)$, as shown in Fig. 6. To derive the PSF for arbitrary motions, we simply need to derive $t(x)$. Fig. 6 shows $t(x)$ for the three types of motion considered here.

6.3 Recovering Motion PSFs

Having shown that many commonly observed motions have closed-form PSFs, we recover B from the amplitude spectrum of a blurred image. Recalling the convolution blur model of (1), the image's amplitude spectrum is

$$|\hat{I}| = |\hat{J}\hat{B}| + |\hat{\eta}| = |\hat{J}||\hat{B}| + |\hat{\eta}|. \tag{8}$$

Since $\hat{\eta}$ has constant expected amplitude at all frequencies, noise is an offset in the frequency domain and can be ignored in our analysis of relative spectral power.

Our goal is to use natural image statistics to recover the PSF from the blurred image I. We assume the motion type (constant velocity, acceleration, etc.) is known and we focus on recovering its corresponding motion parameters α. We can apply the R_u operator to (8):

$$R_u[|\hat{I}|] = \sum_{v=0}^{V} |\hat{J}(u,v)||\hat{B}(u)| = |\hat{B}(u)| \cdot R_u[|\hat{J}|]. \tag{9}$$

Equation (9) allows us to separate $R_u[|\hat{J}|]$ and $|\hat{B}|$. We can further take the log of (9) as:

$$\log(R_u[|\hat{I}|]) = \log(|\hat{B}|) + \log(R_u[|\hat{J}|]). \tag{10}$$

6.4 Motion Estimation Algorithm

Fig. 7 illustrates our motion estimation algorithm. We first determine the motion direction and align it with the u-axis. For every candidate motion parameter α, we compute its PSF B^α and MTF $|\hat{B}^\alpha|$, and estimate the latent image amplitude spectrum $|\hat{J}^\alpha| = |\hat{I}|/|\hat{B}^\alpha|$. We then compute the linear statistics $R_u[|\hat{J}^\alpha|]$ and $R_u[|\hat{I}|]$. Finally, we compute the match score μ between $\log(|\hat{B}^\alpha|)$ and $\log(R_u[|\hat{I}|]) - \log(R_u[|\hat{J}^\alpha|])$. Our estimated motion parameter α is that which maximizes the match score μ.

Estimating the motion direction. Similarly to [22], we find the direction with most muted high frequencies by inspecting the Radon-power spectrum of I in all directions and choosing the one with the maximal variance. This assumes that the latent sharp image is roughly isotropic, i.e., the power spectrum has similar characteristics (variance, mean values) in all orientations.

Computing linear statistics of $|\hat{J}^\alpha|$. A crucial step in our motion estimation algorithm is to derive the linear statistics of $|\hat{J}^\alpha| = |\hat{I}|/|\hat{B}^\alpha|$ from the circular statistics. When J^α is motion blur free, its circular statistics should follow $1/\omega$-exponent distribution. To estimate C and m, we compute the discrete circular averaged power spectrum and apply line fitting between $\log(circ_\omega[|\hat{J}^\alpha|])$ and $\log(\omega)$. We then approximate the linear statistics $R_u[|\hat{J}^\alpha|]$ using (4).

4. Alternately, the method can be applied to motion-blurred regions that have been segmented from the background, as in [2].

Fig. 5. Comparison of poorly reconstructed coded exposure and traditional images. (Left) Estimated latent image for a 4 percent underestimate of coded exposure blur, with an RMSE of 12.7. (Center) Estimate for a 4 percent underestimate of traditional motion blur, with RMSE 12.3. (Right) Ground truth. Despite similar RMSEs, there are sharpness and noise differences between the estimates.

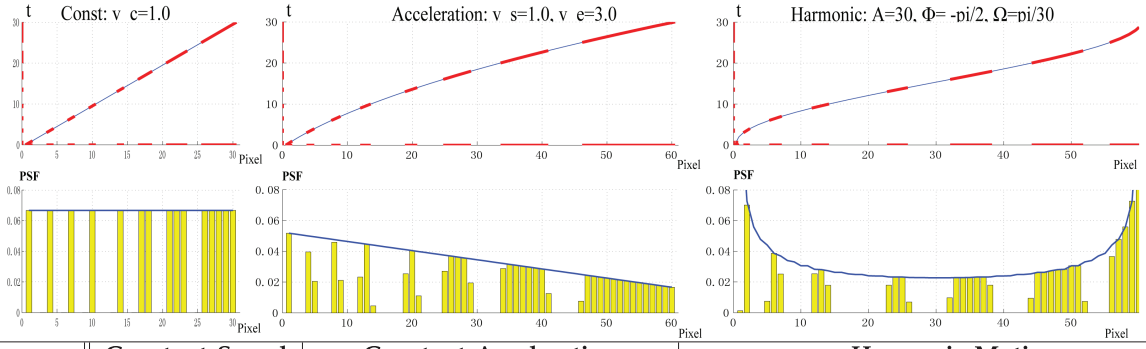

	Constant Speed	Constant Acceleration	Harmonic Motion
Parameters	velocity ν_c	start velocity ν_s, acceleration a	amplitude A, angular speed Ω, initial phase Φ
$t(x)$	$\frac{x}{\nu_c}$	$\frac{-\nu_s+\sqrt{\nu_s^2+2a\cdot x}}{a}$	$\frac{\arcsin(\frac{x}{A}-\Phi)}{\Omega}$
$B_0(x)$	$\frac{S\left(\frac{x}{\nu_c}\right)}{\nu_c}$	$\frac{S\left(\frac{-\nu_s+\sqrt{\nu_s^2+2a\cdot x}}{a}\right)}{\sqrt{\nu_s^2+2a\cdot x}}$	$\frac{S\left(\frac{\arcsin(\frac{x}{A}-\Phi)}{\Omega}\right)}{A\Omega\cos\left(\Omega\cdot\frac{\arcsin(\frac{x}{A}-\Phi)}{\Omega}+\Phi\right)}$

Fig. 6. PSFs of common motions: constant velocity (left), constant acceleration (middle), and harmonic motion (right). The top row shows the time-velocity function sampled by the shutter (in red). The bottom row of figures shows the corresponding PSF.

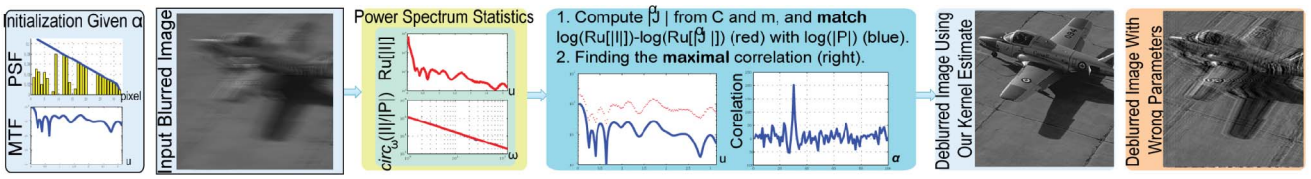

Fig. 7. Steps of our statistical PSF estimation algorithm.

Matching log-linear statistics. Recall that our ultimate goal is to match $f_1 = log(|\hat{B}^\alpha|)$ and $f_2 = \log(R_u[|\hat{I}|]) - \log(R_u[|\hat{J}^\alpha|])$ under a metric μ. A naive μ is to measure the squared difference at sampled points on f_1 and f_2. Since the power spectra of images generally have much smaller values in high frequency, directly computing the correlation between the estimate f_1 and f_2 results in unequal contributions from different frequencies.

Instead, we employ a metric based on the signs of the function derivatives to equally treat all frequencies. Specifically, we use a derivative sign function $\Gamma(\cdot)$:

$$\Gamma(f(u)) = \begin{cases} 1, & \frac{df}{du} \geq 0 \\ -1, & \frac{df}{du} < 0, \end{cases} \tag{11}$$

where f is a 1D function on u. We sample f_1 and f_2 at discrete points $u_1, u_2, \ldots, u_n$, and compute

$$\mu(f_1, f_2) = \sum_{i=1}^{n} \Gamma\big(f_1(u_i)\big)\Gamma\big(f_2(u_i)\big). \tag{12}$$

This functional measures the agreement between the positions of inflection points in f_1 and f_2.

7 PSF Estimation Results

We have applied our technique to the images from [1] and get high-quality reconstructions from estimates within 1 pixel of ground truth. In order to test other motions (acceleration, harmonic) handled by our method, we have acquired additional test images using a Point Gray Flea2 camera. To deconvolve our estimated B, we use the linear system solution [1] for constant velocity motions and the Gaussian-derivative-prior method [23] for constant acceleration and harmonic rotation motions.

Constant velocity. Fig. 8 shows iris and bar code images captured under constant velocity. The motion is from left to right, giving the flutter shutter images Figs. 8a and 8d. We measure the ground truth using a step edge target and find its blurred width in an image with known exposure time. We estimate motion direction within 1 degree of ground truth, and the matching metric μ has

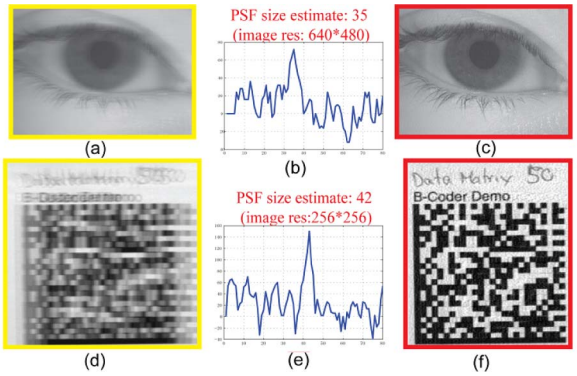

Fig. 8. Results for constant velocity. Column 1: Input images. Column 2: Matching metric versus velocity. Column 3: Deblurred results using our estimated velocity.

maxima at exactly the ground truth values (35 and 42 pixels, respectively). Our latent image estimates Figs. 8c and 8f contain sufficient detail to perform recognition: An iris template extracted from Fig. 8c was successfully matched to another from the same eye, and the barcode in Fig. 8f was successfully decoded. Neither iris recognition nor barcode decoding were successful on Lucy-Richardson [24] deblurred traditional images captured with the same setup.

Constant acceleration. We capture accelerated motion images using a toy car on an inclined track, where gravity provides acceleration, and validate our motion estimation by the quality of the deblurred results. Fig. 9a shows the coded exposure image. Though the track is slanted at 55 degrees, the camera is rotated so that the motion appears nearly horizontal. We first apply our motion direction estimation algorithm, which produces an estimate of 1 degree. Next, we determine the motion parameters, shown in Fig. 9d. The deblurred result (Fig. 9f) and close-up (Fig. 9h) show a significant amount of detail relative to the initial blur. Readers should note that the shutter sequence used was designed to produce an invertible PSF under constant velocity motion, so invertibility with accelerated motion is not assured. Artifacts in Fig. 9f are due to the stationary background's intensity interacting with the moving foreground object, and may be reduced by applying motion segmentation. For comparison, we present the deblurred result assuming constant velocity. As shown in Figs. 9e and 9g, using the incorrect motion model results in severe artifacts.

Harmonic rotation. As shown in Fig. 10a, we emulate harmonic rotation by hanging a rigid object using two approximately rigid, weightless cords connected to the same point. By swinging the object back and forth freely within a plane, we synthesize a periodic harmonic rotation. Note that the rotation gives spatially varying blur kernels and that, like the previous example, the shutter sequence will not produce an invertible PSF for all velocities. We transform the harmonic rotation into a linear motion by tracking feature points and estimating the rotation center using least squares [12]. We then warp the image along the radial directions to form a spatially invariant blur Fig. 10d. Our algorithm recovers the harmonic motion parameters Fig. 10f and then deblurs the image Fig. 10d to obtain Fig. 10e. Last, we unwarp the deblurred image, giving the result in Fig. 10h.

8 Conclusion

We have demonstrated statistical blur estimation providing accurate motion estimates for constant velocity, constant acceleration, and harmonic rotation in real images. By matching inflections in the image's amplitude spectrum to modeled PSFs under different motion parameters, we avoid assuming that the latent image contains a strong edge separating foreground and background. However, our blur estimation method must be applied to a region with a single motion, either a uniformly blurred image or a region segmented by other means.

We have also improved the existing method for determining an optimal coded exposure PSF. We show that the existing criteria which maximizes the MTF minimum does not provide latent image estimates with optimal RMSE. Using natural image statistics, we demonstrate better estimation of the optimal PSF (in the sense of RMSE) in the same amount of computing time. Though not the focus of this work, we point out that this gain results in more than just aesthetic improvements as coded exposure images have been demonstrated to be more useful for both iris biometrics [25] and barcode recognition [26]. While our PSF selection uses RMSE as the objective function, we can substitute application-specific objectives (e.g., equal error rate for iris recognition) for RMSE to address application-specific optimization.

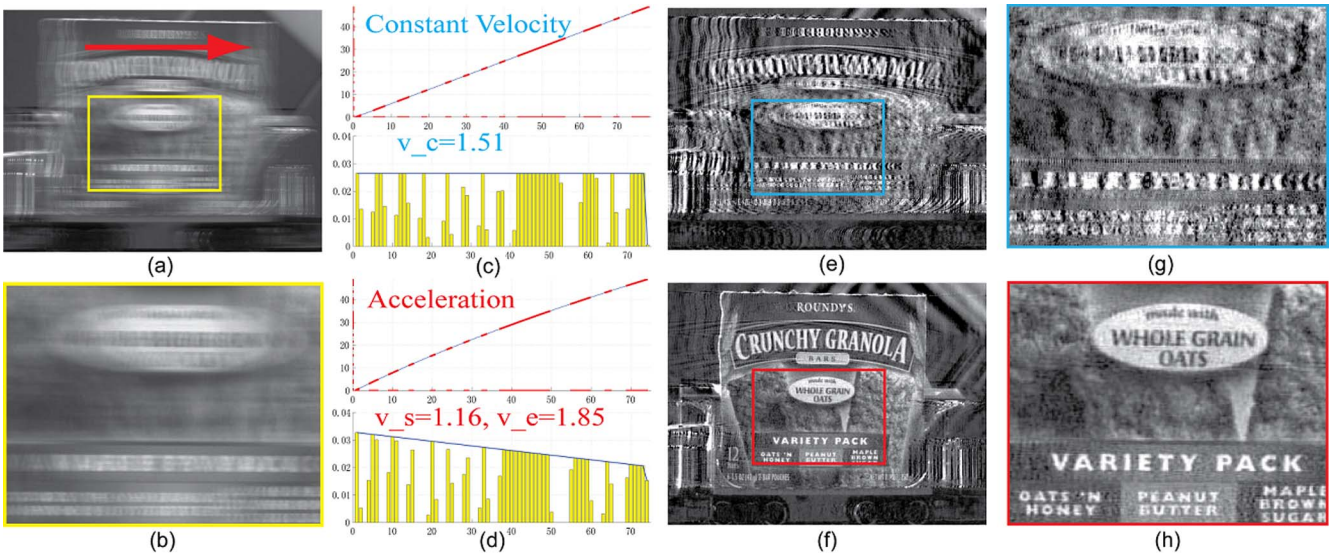

Fig. 9. Deblurring a toy car accelerating down a 55 degree incline. The camera is rotated, giving horizontal motion, and a target is attached to the car. (a) Input image. (c)/(d) $t(x)$ and corresponding PSFs for constant velocity/acceleration. (e) Deblurred result assuming constant velocity PSF. (f) Deblurred result using our estimated acceleration. (b)/(g)/(h) are close-ups of (a)/(e)/(f).

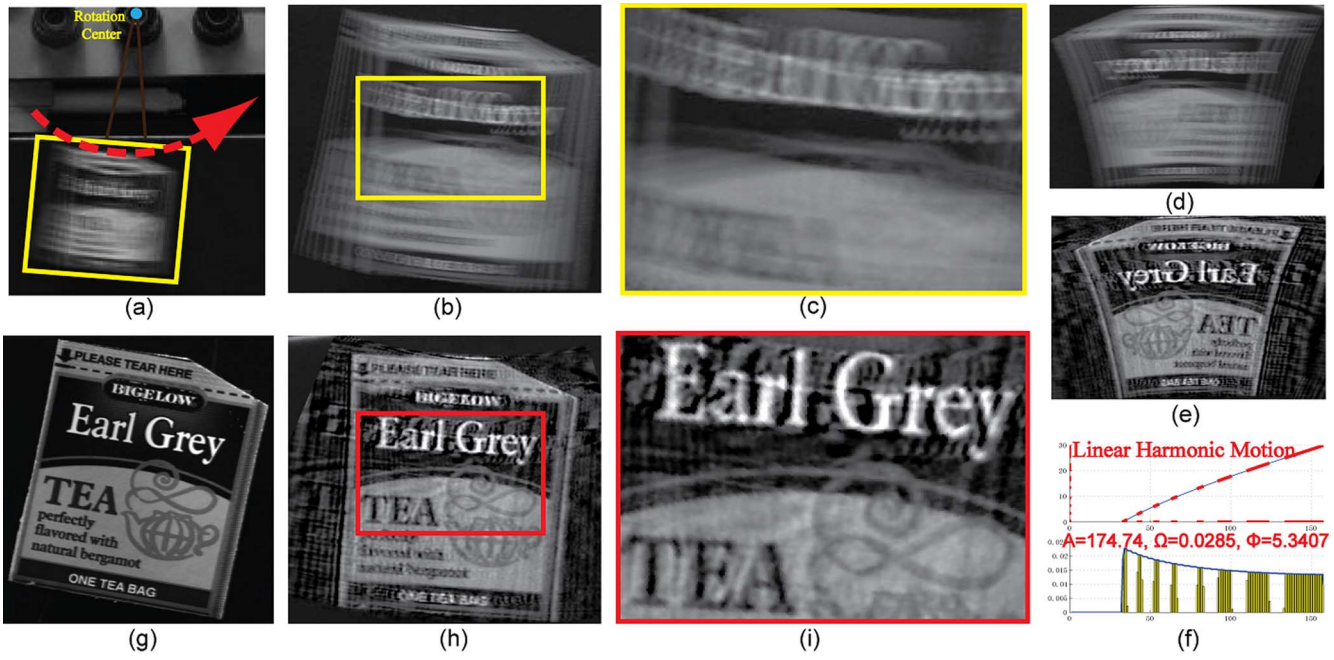

Fig. 10. Results on harmonic rotation (reference image in (g)). (a)-(c) The motion-blurred flutter shutter image. (d) The warped (b) under polar coordinates. (e) and (h) The warped and unwarped version of our deblurred result. (i) A close-up view of the deblurred result. (f) $t(x)$ and our recovered PSF.

ACKNOWLEDGMENTS

This material is based upon work supported by the US Army Biometrics Task Force under Contract No. W91CRB-09-C-0013. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the US Army. Flea2 is a trademark of Point Grey Research, Inc. All brand names and trademarks used herein are for descriptive purposes only and are the property of their respective owners. Drs. Ding and Yu were partially supported by US National Science Foundation grant IIS-CA-REER-0845268 and US Air Force Office of Scientific Research grant FA9550-10-1-0175.

REFERENCES

[1] R. Raskar, A. Agrawal, and J. Tumblin, "Coded Exposure Photography: Motion Deblurring Using Fluttered Shutter," *ACM Trans. Graphics*, vol. 25, no. 3, pp. 795-804, 2006.
[2] A. Levin, "Blind Motion Deblurring Using Image Statistics," *Proc. Neural Information Processing Systems*, pp. 841-848, 2006.
[3] A. Agrawal and Y. Xu, "Coded Exposure Deblurring: Optimized Codes for PSF Estimation and Invertibility," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2009.
[4] C. Zhou and S.K. Nayar, "Computational Cameras: Convergence of Optics and Processing," *IEEE Trans. Image Processing*, vol. 20, no. 12, pp. 3322-3340, Dec. 2011.
[5] S. Haykin, *Blind Deconvolution.* Prentice-Hall 1994.
[6] O. Whyte, J. Sivic, A. Zisserman, and J. Ponce, "Non-Uniform Deblurring for Shaken Images," *Int'l J. Computer Vision*, vol. 98, no. 2, pp. 168-186, 2012.
[7] A. Gupta, N. Joshi, L. Zitnick, M. Cohen, and B. Curless, "Single Image Deblurring Using Motion Density Functions," *Proc. European Conf. Computer Vision*, 2010.
[8] J. Jia, "Single Image Motion Deblurring Using Transparency," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2007.
[9] Q. Shan, J. Jia, and A. Agarwala, "High-Quality Motion Deblurring from a Single Image," *Proc. ACM Siggraph*, 2008.
[10] R. Fergus, B. Singh, A. Hertzmann, S.T. Roweis, and W.T. Freeman, "Removing Camera Shake from a Single Photograph," *ACM Trans. Graphics*, vol. 25, no. 3, pp. 787-794, 2006.
[11] D. Krishnan and R. Fergus, "Fast Image Deconvolution Using Hyper-Laplacian Priors," *Proc. Neural Information Processing Systems*, 2009.
[12] Q. Shan, W. Xiong, and J. Jia, "Rotational Motion Deblurring of a Rigid Object from a Single Image," *Proc. 11th IEEE Int'l Conf. Computer Vision*, 2007.
[13] P. Favaro and S. Soatto, "A Geometric Approach to Shape from Defocus," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 27, no. 3, pp. 406-417, Mar. 2005.
[14] N. Joshi, C. Zitnick, R. Szeliski, and D. Kriegman, "Image Deblurring and Denoising Using Color Priors," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2009.
[15] L. Xu and J. Jia, "Two-Phase Kernel Estimation for Robust Motion Deblurring," *Proc. European Conf. Computer Vision*, 2010.
[16] S. Cho and S. Lee, "Fast Motion Deblurring," *Proc. ACM Siggraph Asia*, 2009.
[17] M.E. Moghaddam and M. Jamzad, "Motion Blur Identification in Noisy Images Using Mathematical Models and Statistical Measures," *Pattern Recognition*, vol. 40, pp. 1946-1957, July 2007.
[18] H. Ji and C. Liu, "Motion Blur Identification from Image Gradients," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2008.
[19] S. Dai and Y. Wu, "Motion from Blur," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2008.
[20] A. van der Schaaf and J.H. van Hateren, "Modelling the Power Spectra of Natural Images: Statistics and Information," *Vision Research*, vol. 36, pp. 2759-2770, 1996.
[21] D. Martin, C. Fowlkes, D. Tal, and J. Malik, "A Database of Human Segmented Natural Images and Its Application to Evaluating Segmentation Algorithms and Measuring Ecological Statistics," *Proc. IEEE Int'l Conf. Computer Vision*, 2001.
[22] J.a.P. Oliveira, M.A. Figueiredo, and J.M. Bioucas-Dias, "Blind Estimation of Motion Blur Parameters for Image Deconvolution," *Proc. Iberian Conf. Pattern Recognition and Image Analysis*, 2007.
[23] A. Levin, R. Fergus, R. Fergus, F. Durand, and W.T. Freeman, "Image and Depth from a Conventional Camera with a Coded Aperture," *Proc. ACM Siggraph*, 2007.
[24] L.B. Lucy, "An Iterative Technique for the Rectification of Observed Distributions," *Astronomical J.*, vol. 79, pp. 745-765, June 1974.
[25] S. McCloskey, W. Au, and J. Jelinek, "Iris Capture from Moving Subjects Using a Fluttering Shutter," *Proc. Fourth IEEE Conf. Biometrics: Theory, Applications, and Systems*, 2010.
[26] W. Xu and S. McCloskey, "2D Barcode Localization and Motion Deblurring Using a Flutter Shutter Camera," *Proc. Workshop Applications of Computer Vision*, 2011.