BlurBurst: Removing Blur Due to Camera Shake using Multiple Images

ATSUSHI ITO
Sony Corporation
and
ASWIN C. SANKARANARAYANAN
Carnegie Mellon University
and
ASHOK VEERARAGHAVAN and RICHARD G. BARANIUK
Rice University

Image deblurring has matured over the last decade; today, there are a wide range of deblurring algorithms that operate successfully in the wild. Yet, there are many applications — including telephoto and low-light photography — where camera shake produces a blur kernel that is large enough to cripple state-of-the-art deblurring algorithms. This failure can be attributed to the decreasing SNR at the higher-frequencies of the latent image with increasing blur kernel size. As a consequence, resolving the finest details in the image is often impossible without undesirable artifacts due to noise amplification. In this paper, we demonstrate that these challenges can be overcome by obtaining multiple blurred images. We make the following observations. First, the burst mode in most digital cameras supports the ability to take a sequence of shots in rapid succession. Second, blur due to camera shake is largely one-dimensional; hence, just obtaining a few blurry images opportunistically produces blur orientations that are not aligned with each other; this produces dramatic improvements in deblurring. Third, an alternating sequence of convex programs can be used to recover both the latent image and blur kernels effectively. We refer to this multi-image deblurring algorithm as *BlurBurst*. We demonstrate applications of BlurBurst in telephoto and low-light photography and highlight broader uses in hand-held high dynamic-range (HDR) imaging.

Categories and Subject Descriptors:

General Terms: Multi-image deblurring

Additional Key Words and Phrases: Deblurring, Multiple images, Telephoto imaging, Low-light photography, HDR imaging

ACM Reference Format:

Ito, A., Sankaranarayanan, A. C., Veeraraghavan, A., Baraniuk, R. G., Bursting Blur: Removing Blur Due to Camera Shake using Multiple Images. ACM Trans. Graph. Submitted

Authors' addresses: land and/or email addresses.

DOI 10.1145/XXXXXXX.YYYYYYY
http://doi.acm.org/10.1145/XXXXXXX.YYYYYYY

1. INTRODUCTION

One of the striking successes of computational photography has been in deblurring images. Image deblurring has traditionally been considered a severely undetermined problem in image processing and computer vision. Yet, today not only do there exist many robust algorithms for solving the problem, but there are also multiple commercial products that successfully operate in the wild.

Behind the success of these solutions lie two fundamental observations. First, blurring is inherently a lossy process. Blur due to camera shake typically introduces nulls in the Fourier spectrum of the images. As a consequence, even when the point spread function (PSF) of the blur, or equivalently the blur kernel, is known, traditional inverse or recovery algorithms based on least squares recovery can fail catastrophically. Second, the loss in information due to the blurring process can often be undone using signal priors. While there are a myriad of image priors used in today's deblurring algorithms, the consistent theme is that of regularization of the inverse algorithm using such priors.

In spite of the tremendous progress in image deblurring, several challenges remain unsolved. The first challenge occurs in telephoto and low-light imaging where the effects of camera shake manifests itself as large blur kernels. Unfortunately, the performance of state-of-the-art deblurring algorithms degrade significantly as the size of the blur kernel increases (see Figure 1(b)). The second challenge is in deblurring scenes that are noisy or contain saturated regions, both of which occur in low-light scenes. While there has been research addressing denoising low-light imagery (cf. Chatterjee et al. [2011]) and telephoto imaging (cf. Joshi and Cohen [2010]), their treatment does not consider blur due to camera shake.

Solving deblurring under high noise and large motion blur requires us to well-condition an otherwise severely ill-conditioned system. In this paper, we propose to obtain multiple blurred images of the same latent image; in this setting, we show that the drawbacks associated with deblurring in high noise and large bluring kernels can be overcome. Almost all modern digital cameras, including cellphones, point-and-shoots, and SLRs, allow the user to capture images in a burst mode wherein a series of images can be quickly acquired. Obtaining multiple images significantly improves

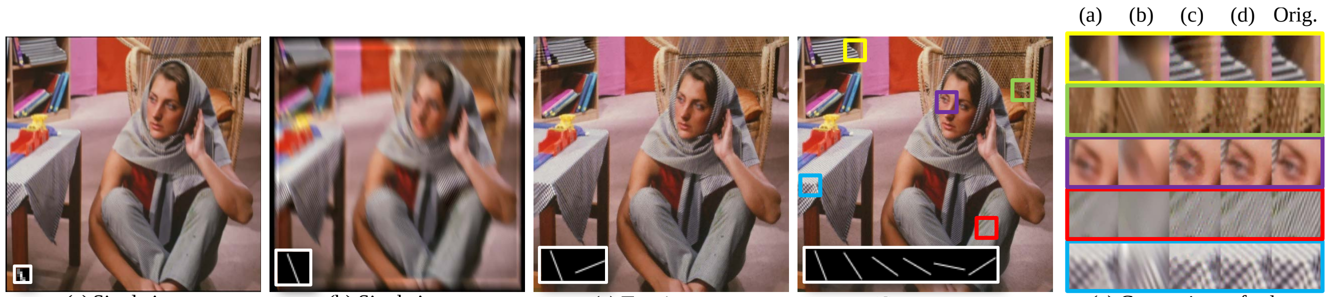

(a) Single image Blur PSF 5 pixels | (b) Single image Blur PSF 25 pixels | (c) Two images Blur PSF 25 pixels | (d) Six image Blur PSF 25 pixels | (e) Comparison of select patches from (a-d)

Fig. 1. Deblurring with large blur kernel sizes. (a) Single image deblurring works well for small blur kernel sizes but (b) fails when the kernel size is increased. (c, d) In this paper, we demonstrate that multi-image deblurring provides dramatic improvement in deblurring even for very large blur kernel sizes. (e) A comparison of the results in (a-d) along with ground truth. Note the dramatic improvement even at two images. Results shown are in the non-blind setting. Inset in (a-d) are the kernels used to generate blurry images.

the latent image recovery in two significant ways. First, obtaining multiple blurred image improves SNR just by virtue of noise suppression. Second, in the case of camera shake, blur kernels tend to be localized close to one-dimensional (1D) curves; hence, the blur kernels have nulls in different locations as well as decay along different orientations while preserving high-frequency information along perpendicular directions. Together, these observations are a game changer. Indeed, just *two* carefully designed blur kernels — each orthogonal to the other — is sufficient for robust image deblurring even when the individual blurs are large (see Figure 1).

Contributions: In this paper, we develop a new methodology called *BlurBurst* that recovers a sharp latent image from multiple blurry input images. We assume that the input images are formed via a single latent image blurred with each different PSFs. Under this assumption, we derive an iterative estimation algorithm that recursively estimates both the unknown blur kernels and the latent image (Section 3). We empirically demonstrate that obtaining multiple blurry images significantly improves recovery performance; specifically, when compared to existing methods with single image input, the proposed method can recover scenes for far greater amount of blur and measurement noise (Section 4). In addition to a suite of simulations, we demonstrate applicability in several real world applications involving large blurs, including telephoto and low-light imaging (Section 5). Finally, we extend BlurBurst to handle deblurring in the presence of saturation to obtain hand-held high dynamic range (HDR) images (Section 6).

2. PRIOR WORK

Deblurring algorithms can be categorized into three groups: single image-based algorithms, multiple image-based algorithms, and methods that rely on special hardware to alter the blurring process in a favorable way.

Single image deblurring: The problem of recovering the latent image from a single blurred image can be reduced to image deconvolution if the blur kernel[1] is shift-invariant. Here, we can sub-divide image deconvolution into non-blind (known kernel) and blind (unknown kernel) cases.

In non-blind deconvolution, a general trend has been to use image priors to regularize the inverse problem of estimating the latent image given blurred image and kernel. Richardson-Lucy (RL) deconvolution [Lucy 1974][Richardson 1972], a technique originally proposed in the 1970s, models pixel intensities of the latent image as Poisson distributed and derives a computational efficient algorithm for this specific image model. There are many other priors that have successfully been used for non-blind deconvolution, including sparse gradients [Levin et al. 2007], sparse wavelet priors [Neelamani et al. 2004], minimum total variation [Osher et al. 2005], and bilateral edge regularization [Yuan et al. 2008].

There have been many methods for blind deblurring from single image toward recovering both a sharp image and the blurring kernel. Seminal work by Fergus et al. [2006] demonstrated that it was indeed possible to deblur real world images reliably; in particular, under a sparse kernel prior and a mixture-of-Gaussian prior on the image gradients, they demonstrated that the expected mean of the kernel conditioned on the observed blurry image serves as a good estimate for the unknown blur kernel. Non-blind deblurring using RL deconvolution is used to obtain the latent image. Cho et al. [2011] make the observation that shape-profile of a blurred edge encodes the shape of blur-kernel perpendicular to the edge orientation. Exploiting this, they use multiple blurred edges to reconstruct the blur kernel under an inverse Radon transform framework. In contrast, Shan et al. [2008] use an iterative algorithm that performs alternating optimization of blur kernel and latent image till convergence; a hallmark of their method is the use higher-order models on the spatial distribution of noise to provide highly accurate blur kernel and latent image estimates. Cho and Lee [2009] introduce inverse-diffusion shock filter to reconstruct sharp edges and fast Fourier transform to recover the blur kernel in gradient domain.

Multiple image deblurring: There are relatively fewer algorithms devoted to multi-image deblurring. The idea of using two images of varying exposures has been explored in Yuan et al. [2007]; the short exposure image, although noisy, has little blur and is used to as a guide for deblurring the less-noisy, but blurry long-exposure image. Agrawal et al. [2009] exploit the idea that blur kernels of different sizes have nulls at different locations in the frequency spectrum. They use this in the context of motion deblurring of object moving in a straight line. Rav-Acha and Peleg [2005] exploit directional properties of camera shake blur and show thatimages with different blurring directions can be used for estimating blur kernel. Liu and Abbas [2003] rely on a high-speed camera to capture multiple im-

[1] We use the terms "kernel" and "PSF" inter-changeably in this paper. Also, we define the "size" of the kernel as the width of the smallest square that encloses the blur kernel.

ages of varying exposures which are fused to recover a motion-free HDR image; here, the blurred regions in the longer exposures are replaced with their corresponding regions from the shorter exposure images. Harmeling et al. [2010] use a generalized expectation maximization framework to propose a single sharp image from a sequence of blurry images; however, the treatment here is only for the case of astronomical images.

The closest to the work presented in this paper are the multi-image deblurring results by Sroubek and Flusser [2003][2005] and Sroubek and Milanfar [2012]. In particular, we share many of the modeling and optimization techniques proposed in Sroubek and Milanfar [2012]. The key differences are the use of a higher-order noise model and our focus on pre-registration of the blurry images, both of which enable us to process a larger set of photographs captured under real-acquisition conditions. Further, in many ways, our optimization framework is significantly simpler, in that we do not use robust statistics. In-spite of this, our proposed algorithm outperforms this algorithm on all datasets that we tested on.

Hardware designs for deblurring: A powerful method for image deblurring is the use of computational optics for actively shaping the blur kernel; this dramatically reduces problems associated with traditional deblurring such as nulls in the frequency spectrum and low SNR at higher frequencies. Ben-Ezra and Nayar [2003] propose a hybrid camera that uses a combination of a high temporal (but low spatial) resolution camera and a high spatial (but low temporal) resolution camera. Motion estimates obtained from the former is used to deblur the high resolution frame from the latter. Raskar et al. [2006] introduce shutter-coding (or the "shutter flutter") for the problem of deblurring 1D motion; coding the shutter shapes the Fourier spectrum of the blur kernel to be as flat while maximizing the SNR. As a consequence, the linear system associated with the blurring process is well-conditioned and invertible. Joshi et al. [2010] use a combination of gyroscopes and accelerometers in order to estimate a blur function from the camera's acceleration and angular velocity during exposure. Levin et al. [2008] construct a motion invariant blur by introducing camera motion; specifically; when the camera is moved on a line with a parabolic displacement profile, the blur kernel associated all objects moving at constant speeds along the same line is invariant to the object speed.

Deblurring in HDR imaging: One of the most widely used technique for HDR imaging is the exposure bracketing scheme [Debevec and Malik 1997], where an HDR image is derived from a series of images with increasing exposures. While this method works exceedingly well when the images are acquired using a tripod, hand-held HDR image acquisition is challenging due to motion blur-induced artifacts in the longer exposures. In addition to the blur, images with longer exposure often contain large regions of saturated pixels; unless dealt with carefully, saturation violates the image formation model underlying most traditional deblurring methods and produces unacceptable artifacts.

There has been some preliminary work devoted to addressing hand-held HDR imaging. Yuan et al. [2008] proposed an approach taking the advantage of having a pair of blurred/noisy images. Lu et al. [2009] produced a unified probabilistic model for estimating blur kernels simultaneously with recovering a HDR irradiance map and camera response curve. Cho et al. [2011] provided detailed analysis on various types of outliers including around saturated areas. They also proposed an EM-based deconvolution method which explicitly detecting and properly handling outliers in the deconvolution process.

3. BLURBURST: A MULTI-IMAGE DEBLURRING ALGORITHM

In this section, we outline our multi-image deblurring algorithm. We refer to this algorithm as *BlurBurst* since we seek to deblur by taking multiple images in rapid succession or in burst mode.

Problem definition: Given a set of Q blurred images $\{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_Q\}$ satisfying the image generation equation:

$$\mathbf{y}_i = \mathbf{k}_i * \mathbf{x} + \omega_i, \quad i = 1, \dots, Q, \tag{1}$$

our goal is to estimate the latent image $\mathbf{x}$ and the blur kernels $\{\mathbf{k}_1, \dots, \mathbf{k}_Q\}$ under a Gaussian noise model on the measurement noise ω_i. Each blurred image $\mathbf{y}_i$ is assumed to be of the same size $M \times M$, pixels and each blur kernel is of size $K \times K$.

Overview of the solution: We employ the following signal priors in BlurBurst: a sparse image gradient prior for the latent image $\mathbf{x}$ (obtained using a minimum total-variation (TV) norm) and a sparse prior on the kernels $\mathbf{k}_i$ (obtained using a minimum ℓ_1 norm). Under these priors, the deblurring problem can be reduced to solving the following optimization problem:

$$\arg\min_{\mathbf{x}, \mathbf{k}_i} \sum_i \|A(\mathbf{y}_i - \mathbf{k}_i * \mathbf{x})\|_2^2 + \lambda_1 \mathrm{TV}(\mathbf{x}) + \lambda_2 \sum_i \|\mathbf{k}_i\|_1 \tag{2}$$

where $A(\cdot)$ is a high-order noise model that we discuss below. The optimization problem in (2) is bi-convex. We solve it as a sequence of convex problems — at each step, optimizing over $\mathbf{x}$ or $\mathbf{k}_i$ with the other variables held fixed at the latest estimates.

Boundary: We handle boundary-related artifacts by optimizing over a latent image that is larger in size than the blurred image. Given blur kernels of size $K \times K$ and blurry observations of size $N \times N$, the size of the estimated latent image is set as $(N - M + 1) \times (N - M + 1)$. Convolution of the latent image with a blur kernel as in (1) returns back the central $M \times M$ pixels that correspond to pixels which are completely determined by the latent image without zero-padding or any other boundary assumptions. This way of handling boundary artifacts can be viewed as a free boundary condition and helps remove ringing artifacts.

Higher-order noise model: We use a likelihood model, first introduced in Shan et al. [2008], that enforces the Gaussian measurement noise model up to multiple spatial derivatives. Equation (1) models the noise term $\omega_i = \mathbf{y}_i - \mathbf{k}_i * \mathbf{x}$, as i.i.d. Gaussian. Viewing ω_i as an $N \times N$ image, Shan et al. [2008] make the observation that spatial derivatives of ω_i are Gaussian as well. Given that spatial derivatives are high-pass filters, adding them to the optimization framework preferentially penalizes the errors at higher-frequencies. As a consequence, enforcing this property leads to sharper blur kernel and latent image estimates.

We capture the higher-order noise models via the operator $A(\cdot)$ in (2). We define A as

$$A(\omega) = \left[\omega \quad \frac{\omega_x}{\sqrt{2}} \quad \frac{\omega_y}{\sqrt{2}} \quad \frac{\omega_{xx}}{2} \quad \frac{\omega_{xy}}{2} \quad \frac{\omega_{yy}}{2} \right],$$

where $\omega_x = \frac{\partial}{\partial x}\omega$, and so on. The scaling term associated with each partial derivative is to account for the increase in noise variance upon application of the derivative operator.

Blur kernel estimation: The optimization problem for estimating the individual blur kernels given an estimate of the latent image $\widehat{\mathbf{x}}$ can be reduced to

$$\arg\min_{\mathbf{k}_i} \|A(\mathbf{y}_i - \mathbf{k}_i * \widehat{\mathbf{x}})\|_2^2 + \lambda_2 \|\mathbf{k}_i\|_1.$$

Here, we use SPG-L1 [Van Den Berg and Friedlander 2008], a fast solver for a variant of this problem.[2]

Latent image estimation: Given estimates of the blur kernels $\{\widehat{\mathbf{k}}_i;\ i = 1, \ldots, Q\}$, estimating the latent image reduces to the following convex problem:

$$\arg\min_{\mathbf{x}} \sum_i \|A(\mathbf{y}_i - \widehat{\mathbf{k}}_i * \mathbf{x})\|_2^2 + \lambda_1 \mathrm{TV}(\mathbf{x}).$$

We use M-Fista [Beck and Teboulle 2009] — a fast, second-order solver for minimum TV-norm problems. For computational speed, the convolution operations are performed in the Fourier domain.

Handling color: The multi-image deblurring algorithm is run only on the luminance (or grayscale image). Once it converges, the final blur kernel estimates are used to perform a non-blind deconvolution on the individual color channels to obtain a color latent image.

Pre-processing/Registration: The blurred images typically require some registration before we can deblur them. The need for registration stems from two factors. First, large displacements between the first and last image would require using an exceptionally large blur kernel; this would increase the computational burden of the recovery algorithm significantly. Second, individual blurry images might have small camera rotations between them which would violate the single latent image with spatially invariant blur model assumed in (1). For these reasons, we introduce a simple homography-based registration step before multi-image deblurring.
The registration pipeline in BlurBurst is as follows:

—**Single-image deblurring:** Given the input blurred images, we first use the single image deblurring algorithm of Shan et al. [2008] to obtain individual latent images. In practice, this step is optional. We observe that the next few steps on feature extraction and matching works only slightly worse on the actual blurry inputs; i.e, in practice, the blurred images can be directly registered without the need for single-image deblurring.

—**Feature extraction and matching:** We extract SIFT features from these deblurred images and match them to a pre-selected reference image.

—**Homography estimation:** The feature correspondences obtained from the matching algorithm are used to fit a homography transformation using RANSAC [Fischler and Bolles 1981]. RANSAC makes our algorithm robust to mismatches due to outliers, poor deblurring, or blurry features.

—**Registered blurred images:** The blurred images are registered using the estimated homography parameters to give us the registered blurred images. At this step, we reject any image that has very poor registration with the reference image.

Once we have the registered blurred images, we crop them to eliminate missing information (this is optional. It is a straightforward extension to incorporate missing data in our framework using a mask on the observed blurred images).

Initialization: Recall that our multi-image deblurring algorithm is iterative, alternating between blur kernel estimation and latent image estimation. To kick-start the alternation, we obtain an initial estimate of the latent image obtained by one of two processes. The first initialization strategy is the use of the one of the blurred images as the latent image. This works well in practice and is a simple method to obtain the initial estimate, but typically takes many more iterations to converge. An alternate initialization strategy is to use the output of a single-image deblurring algorithm. We use the single-image deblurring algorithm of Shan et al. [2008] on all the observed images to obtain multiple candidates for the latent image. We pick among these using a image quality metric. The deblurred image with the best quality metric is used to kick-start the kernel estimation step. We do note that this is tightly coupled with the registration process outlined above; typically, the first step of the registration process provides us with the initial estimate as a by product.

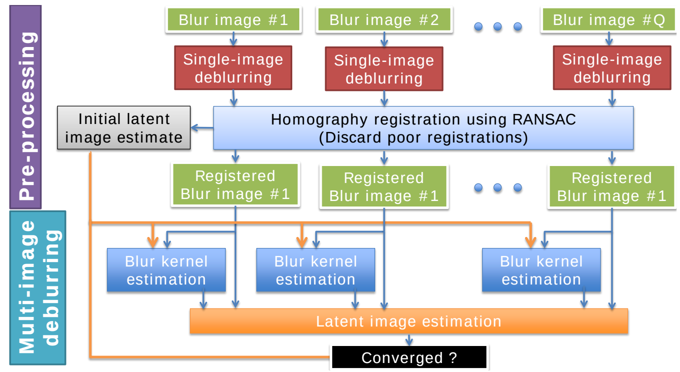

Fig. 2. Outline of BlurBurst, the proposed multi-image deblurring algorithm, including the pre-preprocessing steps and the deblurring algorithm.

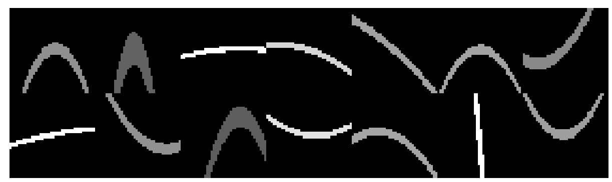

Fig. 3. Example blur kernels simulated as a sub-level set of second-order polynomials with random coefficients over the spatial axes.

The multi-image deblurring algorithm is now applied on the registered blurred images starting with blur kernel estimation. Figure 2 outlines this pipeline.

We do note here that images that do not register well are discarded to avoid artifacts due to model mismatch.

4. ANALYSIS

In this section, we analyze the performance of BlurBurst and contrast it to traditional single-image deblurring. We focus both on a theoretical study, that focuses on “invertibility” of non-blind deblurring in single and multi-image settings as well as an empirical analysis where we compare against state-of-the-art single image algorithms under varying noise, blur size, and number of blurry images.

For the simulations in this section, we randomly generated motion blur kernels as a sub-level set of a conic with random coefficients. Figure 3 shows several examples of blur kernels that were generated using this method. Notice that the blur kernels tend to have 1D features, which effectively simulate camera shake blur kernels.

[2]We use the basis pursuit denoising (BPDN) version of the problem.

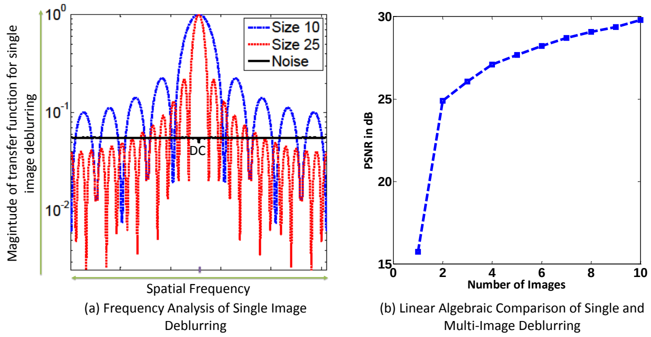

Fig. 4. (a) Frequency Analysis of single image deblurring showing that even in faily small noise level ($\sigma = \frac{1}{255}$) high frequency detail information is completely swamped by observation noise. (b) The mean-squared error of a maximum-likelihood estimator under no signal prior assumptions clearly shows the significant benefits of multiple images over single-image deblurring.

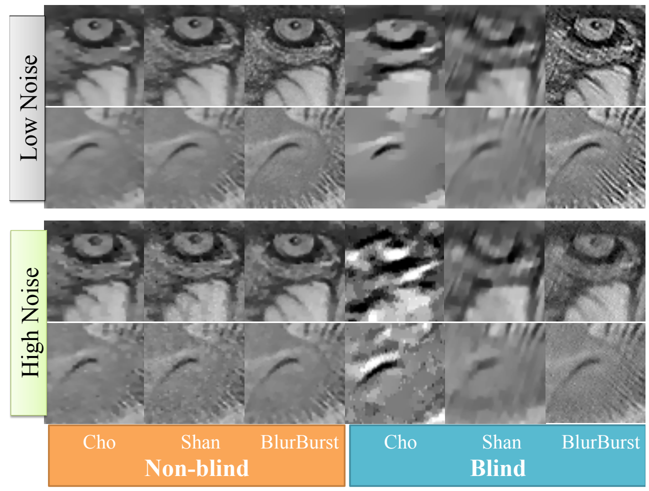

Fig. 5. Comparison of BlurBurst with state-of-the-art single-image deblurring algorithms in low and high noise. Comparisons with both the blind (i.e., psf unknown) and the non-blind (psf known) versions of all algorithms are shown. In both the blind and non-blind cases, BlurBurst outperforms single-image algorithms.

4.1 Theoretical analysis of multi-image deblurring

Pitfalls of single-image deblurring: Blur destroys high frequency information selectively. This means that, in the presence of even minimal observation noise, high frequency details cannot be robustly recovered. This can be clearly visualized by looking at the magnitude of the Fourier transform of the blur kernels and comparing them with the noise floor. Figure 4(a) illustrates a comparison of the frequency spectra of the blur kernels of size 10 and 25 pixels and white Gaussian noise of standard deviation 1/255, which would be considered low noise. In the entire paper, we assume image intensity values in $[0, 1]$. Even under such low noise conditions, notice that the large 25-pixel blur kernel results in a complete loss of high frequency information; this is the primary obstacle that limits the performance of single-image deblurring. Traditional methods to tackle this include PSF engineering to make the blur kernal broadband [Raskar et al. 2006; Levin et al. 2008] and incorporating signal priors to regularize the resulting deblurring problem [Shan et al. 2008; Cho and Lee 2009; Fergus et al. 2006].

Analysis of multi-image deblurring: Here, we analyze the characteristics of multi-image deblurring using a linear algebraic characterization.

The multi-image deblurring problem, as defined in (1), reduces to a over-determined (but possibly ill-conditioned) linear system *when* we have knowledge of the blur kernels. In such a setting, the non-blind deblurring problem can be succinctly written as $\mathbf{y} = \Phi\mathbf{x} + \mathbf{n}$, where $\mathbf{y} = (\mathbf{y}_1, \ldots, \mathbf{y}_Q)$ are the observed blurred images, Φ is the multiplexing matrix encompassing the individual convolutions due to blurring, and $\mathbf{n}$ is i.i.d. Gaussian noise with standard deviation σ. The maximum likelihood (ML) estimate of $\mathbf{x}$, $\widehat{\mathbf{x}}$ is given by

$$\widehat{\mathbf{x}} = (\Phi^T\Phi)^{-1}\Phi^T\mathbf{y}.$$

Thus the covariance matrix Σ of the errors $\mathbf{x} - \widehat{\mathbf{x}}$ in the estimate is given by [Poor 1988]

$$\Sigma = \sigma^2(\Phi^T\Phi)^{-1}\Phi^T\Phi(\Phi^T\Phi) = \sigma^2(\Phi^T\Phi)^{-1}.$$

Therefore, the per-pixel mean-squared error in the estimate is given by

$$\mathrm{MSE} = \mathrm{trace}(\Phi^T\Phi)^{-1}\frac{\sigma^2}{m}, \qquad (3)$$

where m is the size of the observation vector $\mathbf{y}$. Fortunately, since both single-image deblurring and multi-image deblurring can be cast as linear inversion problems with appropriate multiplexing matrices, the MSE performance of both single and multi-image systems can be analyzed (under the assumption of no signal priors) using (3). Figure 4(b) shows the predicted PSNR, i.e., $20\log_{10}(\frac{1}{\mathrm{MSE}})$, for the deblurring problem as a function of the number of images used in deblurring when the blur kernel is of size 25 pixels and the noise standard deviation is $1/255$ just as in Figure 4(a). Here, the blur kernels were simulated using the procedure described earlier. Not surprisingly, single-image deblurring produces very poor reconstructions simply because the high frequency information is completely over-powered by the noise (compare the red and black curves in Figure 4(a)). Even adding *just one* additional blurred image significantly improves the predicted performance, since camera shake blur kernels tend to be nearly 1D[3] and the primary direction of the blur kernel is likely to be at least slightly different in the two blurred images. Additional blurred images continue to provide performance improvements. The analysis presented above assumes that the blur kernels are known and that there are no registration errors. In practice, the analysis still serves as a guide on the achievable performance gains in the best case.

4.2 Performance characterization of multi-image deblurring

In order to study the characteristics of the BlurBurst algorithm and characterize its limitations and performance improvements we perform a series of carefully constructed simulation experiments on a small database of images commonly used in image processing.

Comparison with state-of-the art methods: We compare our results with those the single-image deblurring algorithms of Cho and Lee [2009] and Shan et al. [2008]. In order to make a fair comparison with these single-image algorithms, we adopt the following

[3]Camera shake blur is still 2D. However, it is localized close to 1D curves in the 2D image space.

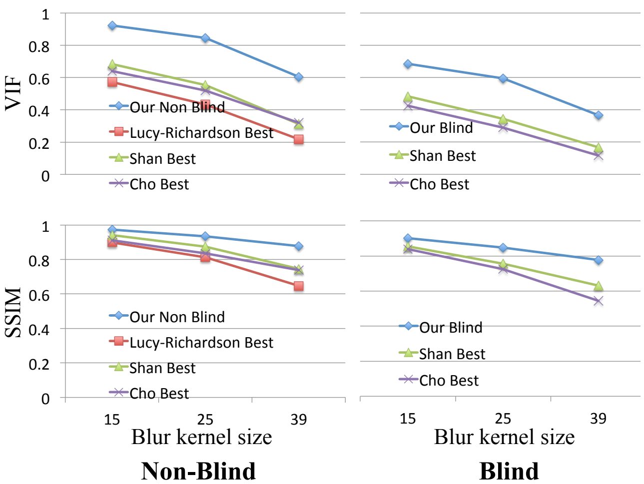

Fig. 6. Performance characterization of deblurring algorithms as a function of the size of blur kernel. As blur kernel size increases performance as measured by two perceptual image quality metrics (VIF and SSIM) decrease. In all cases, BlurBurst performs better than competing single image deblurring algorithms.

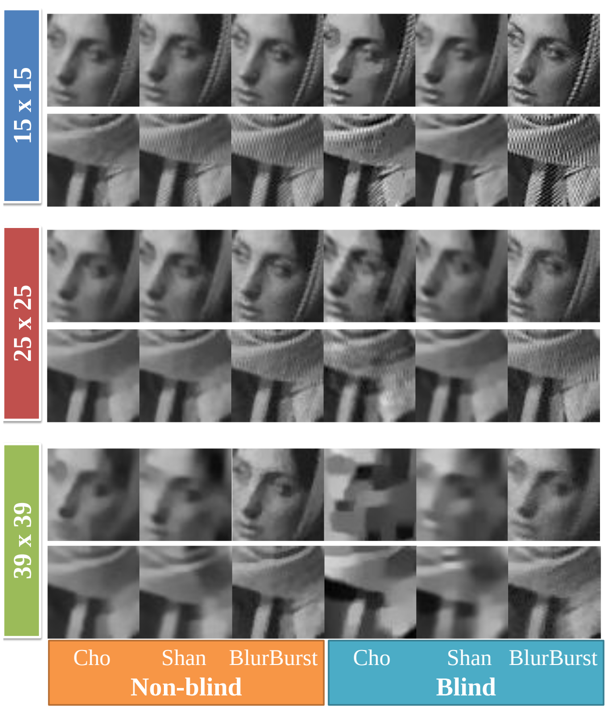

Fig. 7. Performance comparison of deblurring algorithms as a function of blur kernel size on Barbara image. Even at large blur kernel sizes, BlurBurst recovers texture detail lost in other methods.

method. For the single-image algorithms, we perform individual deblurring with each of the captured images. Naturally, the motion characteristics in the blurred images are different leading to significant differences in the amount of blur in the individual images. Therefore, some of the images are deblurred better than the others.

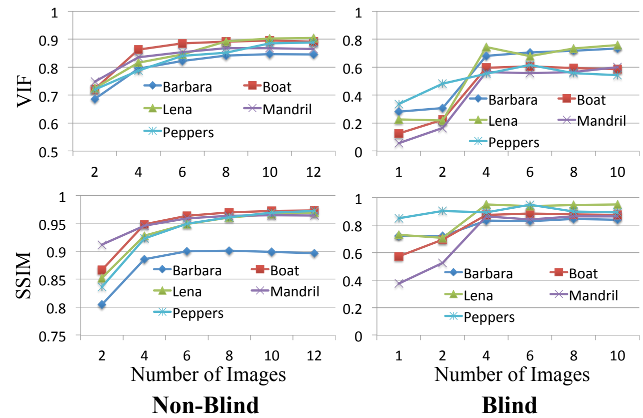

Fig. 8. Performance characterization of BlurBurst as a function of the number of images used. As the number of blurred images used increases performance as measured by two perceptual image quality metrics (VIF and SSIM) increases initially but then saturates indicating that a few images are sufficient to get high quality deblurring results.

We show only the best deblurred image from the set of single-image deblurred images for each set.

Effect of image noise: We compare the performance of blind and non-blind versions of our algorithm with those of Cho and Lee and Shan et al. on a set of commonly used images like Lena, Barbara and Baboon. In this comparison, we used 6 blurred images as input and the blur kernel size was approximately 25 pixels. Figure 5 shows a comparison of the deblurred images. Note that BlurBurst outperforms single image algorithms and recovers texture detail accurately even in high noise scenarios.

Effect of blur kernel size: Figure 6 shows a comparison of deblurring algorithms as a function of the size of blur kernel. We added noise of standard deviation 4/255 to each of the 6 blurry images before performing the deblurring. As blur kernel size increases performance as measured by two perceptual image quality metrics, VIF [Sheikh and Bovik 2006] and SSIM [Wang et al. 2004] decrease. In all cases, our multi-image deblurring algorithm outperforms competing single-image deblurring algorithms, often by a significant margin. A few zoomed in regions from the deblurred images of Barbara are shown in Figure 7. Even at large blur kernel sizes BlurBurst recovers texture detail lost in other methods.

Effect of number of images: In order to estimate the number of images required to significantly improve image deblurring performance, we performed a series of simulations with increasing number of blurry images. Each blur kernel was about 25 pixels wide, and we added noise of standard deviation $4/255$ intensity levels. Figure 8 shows the performance improvement as a function of the number of images being used. Clearly, using two images provides a significant benefit over using a single image. Further, multi-image deblurring performance begins to saturate at about 6 images, indicating that the burst mode in almost all digital still cameras and SLR cameras can be effectively used to capture the multiple images required for multi-image deblurring.

5. REAL IMAGE EXPERIMENTS

We focus on our two motivating applications: telephoto imaging and low-light photography.

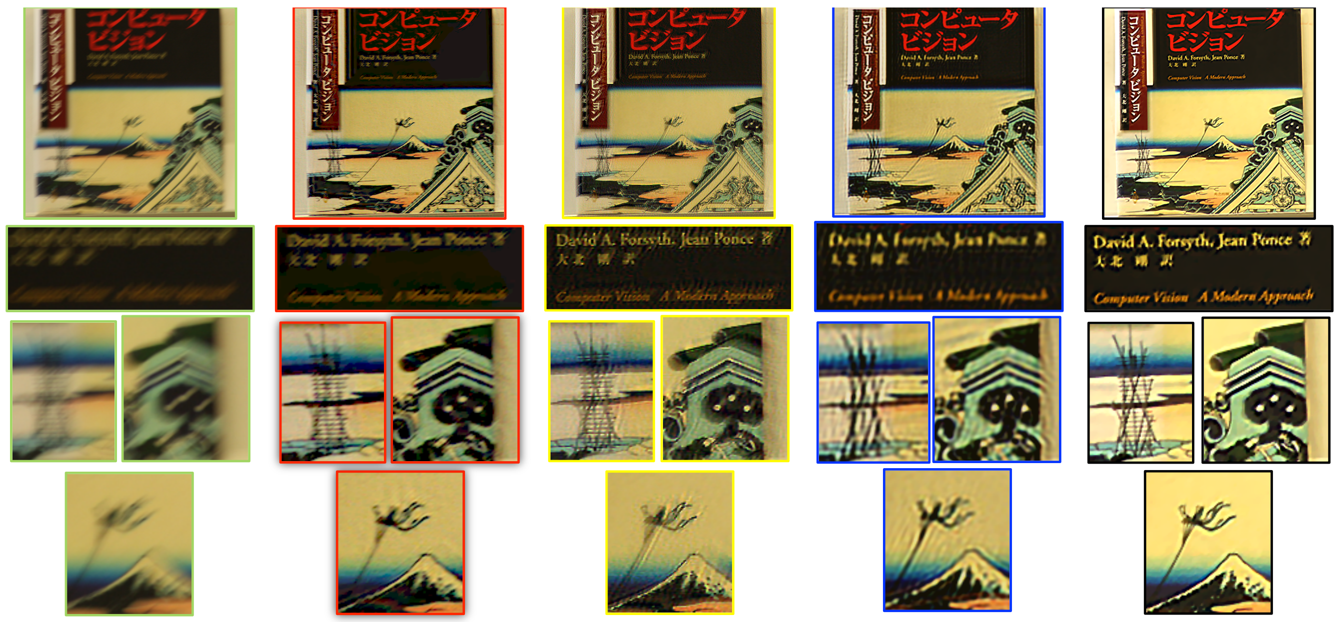

(a) Blurry image (b) Cho and Lee 2009 (c) Shan et al. 2008 (d) Sroubek Milanfar 2012 (e) BlurBurst

Fig. 9. Telephoto imaging with a 300mm lens. Camera shake during telephoto imaging can lead to large amounts of blur. Here, we obtained and processed 6 images at 512×512 resolution. (a) One of the 6 blurry input images. (b-c) Deblurred results using single-image deblurring algorithms of Cho and Lee [2009] and Shan et al. [2008]. In each case, we show the best of the 6 deblurred images. (d) Deblurring results from the state-of-the-art multi-image deblurring algorithm of Sroubek and Milanfar [2012]. (e) Result obtained using BlurBurst. In each column, the top row corresponds to the image and shown below are select patches.

Telephoto imaging: In telephoto imaging, the large zoom implies that each pixel on the camera is associated with a very narrow cone of light. Here, even the slightest camera shake leads to dramatic blurring artifacts. We showcase this with multiple examples. Shown in Figure 9 are deblurring results on a "book" dataset. We collected images of a book using a Canon camera operating at ISO 400, F/8, exposure time 1/20 second, and focal length 300mm. The latent image recovered using BlurBurst captures very fine textures (such as the fence and the tree) faithfully. Figure 10 shows additional details on intermediate estimates, blur kernel estimates as well as individual deblurring results from state-of-the-art single image deblurring algorithms.

Figure 11 shows the performance of BlurBurst and competing algorithms on a images of a resolution chart. The images were collected with the following parameters: ISO 100, F/20, 1/4 sec exposure and focal length 300mm. The result obtained using BlurBurst is striking in its ability to reproduce the finest details on the resolution chart that is as good as that observed in the tripod image. We invite the reader to use the PDF zoom tool to look at smaller details on the individual images in the electronic version.

Figures 12, 13, and 14 encapsulate deblurred results of BlurBurst over many different datasets.

Low-light photography: In low-light photography, the longer-exposures required to obtain high-quality images also introduce large blur due to camera shake. Figure 15 shows an example. The images were captured with a Canon SLR with the following settings: ISO 400, F/5.6, exposure times 1.6 seconds, and focal length 55mm. Note the large blur on the input blurry image; state-of-the-art single image deblurring algorithms introduce artifacts as a consequence. BlurBurst recovers all the subtle details faithfully; specifically, concentrate on the hands and facial features of the baby. Figure 16 shows additional details on intermediate estimates, blur kernel estimates as well as individual deblurring results from state-of-the-art single image deblurring algorithms. Results on a second dataset obtained in night-light is shown in 17; these were obtained with ISO100, F/8, and exposure time 0.25 seconds.

6. HAND-HELD HIGH DYNAMIC-RANGE IMAGING

A simple method to increase the dynamic-range of a camera is to obtain multi-image of varying exposures and fuse them to obtain a HDR image of the scene. Invariably, this involves obtaining images with long exposures which, in the absence of tripods, lead to significant blur due to camera shake. In this section, we extend BlurBurst to hand-held HDR image acquisition.

Challenges: There are two key challenges in deblurring for HDR imaging: first, the presence of saturated pixels; and second, signal-dependent noise due to varying exposures. The presence of saturated pixels has historically been the point of failure for most of the existing deconvolution algorithms. These are caused when the radiance of the scene exceeds the range of the camera's sensor, leaving bright highlights clipped at the maximum output value (e.g. 255 for an 8-bit image). When we shoot a low-light scene with a long exposure time, this effect should be seen in and around a few bright spots. Deblurring images without accounting for the non-linearity induced by saturation leads to severe ringing artifacts. Consider the example of non-blind deblurring in the presence of saturation shown in Figure 18. The deblurred image, obtained using Shan et al. [2008], exhibits heavy ringing artifacts around the saturated regions. This is a consequence of the violation of the forward blur-

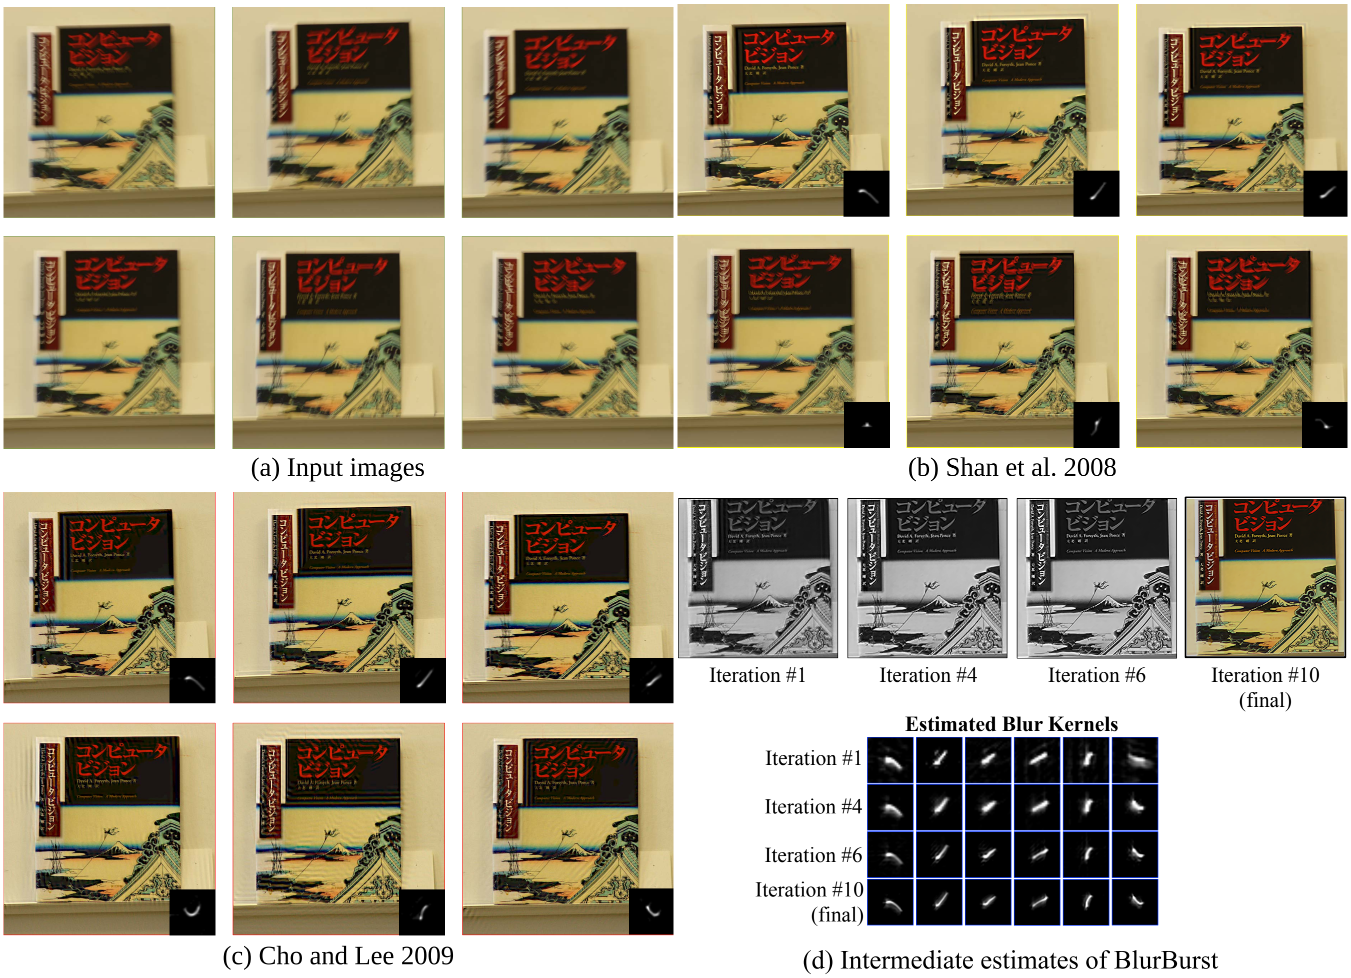

Fig. 10. Telephoto imaging result. Shown are (top-left) input blurry images, results using single-image deblurring algorithms of (top-right) Shan et al. and (bottom-left) Cho and Lee. Finally, (bottom-right) shown are intermediate estimates of latent image and blur kernels from BlurBurst.

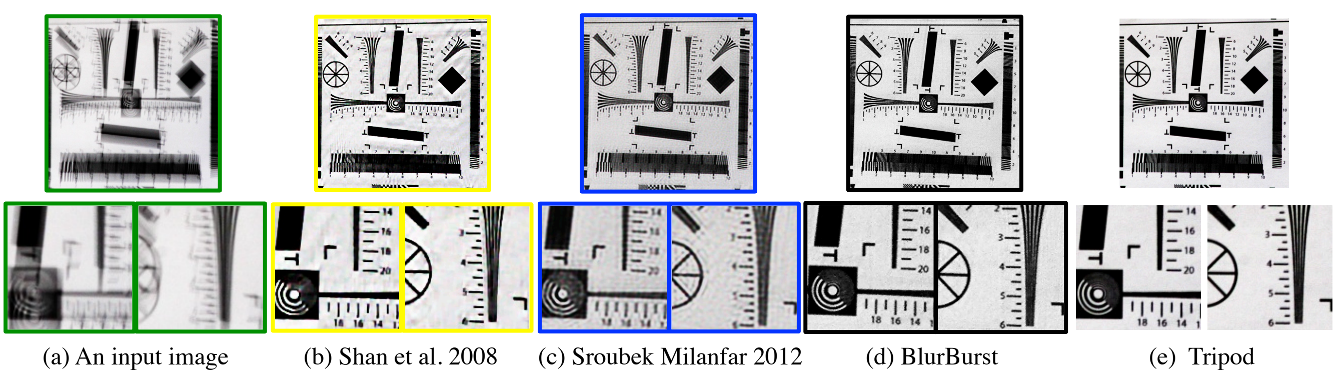

Fig. 11. Telephoto and low-light imaging result of a resolution chart. Shown are (a) an input blurred image, (b-d) deblurring results from various algorithms, and (e) a tripod stablized sharp image. For each result, we show the deblurred image on top and two zoomed-in insets below.

ring model due to saturation and the propagation of error due to model misfit.

Problem definition: Given a set of Q blurred images (corresponding to increasing exposure durations) $\{\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_Q\}$, unsatu-

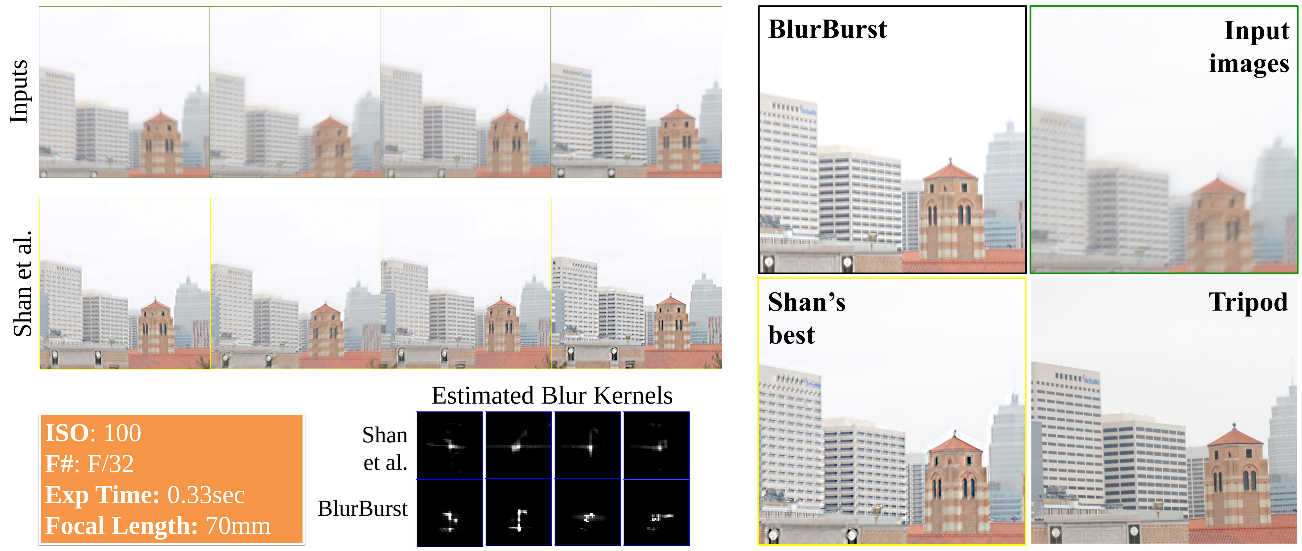

Fig. 12. Telephoto imaging result. (Left) Shown are (top) input blurry images, (middle) Shan et al. results, (bottom) blur kernel estimates of Shan et al. and BlurBurst. The inset in orange shows camera parameters uses. (Right) Zoomed-in versions of the deblurred outputs and a tripod image.

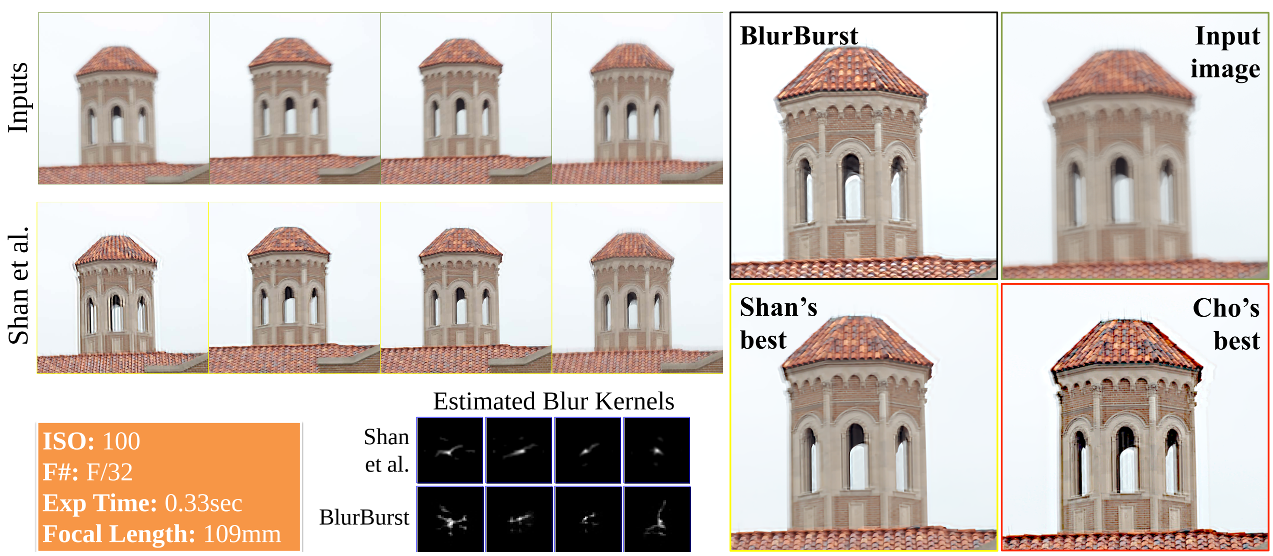

Fig. 13. Telephoto imaging result. (Left) Shown are (top) input blurry images, (middle) Shan et al. results, (bottom) blur kernel estimates of Shan et al. and BlurBurst. The inset in orange shows camera parameters uses. (Right) Zoomed-in versions of the deblurred outputs.

rated pixels satisfy the image formation model given by

$$\mathbf{y}_i^{NS} = (\mathbf{k}_i * \mathbf{x})^{NS} + \omega_i, \quad i = 1, \ldots, Q,$$

where $(.)^{NS}$ represents the set of all pixels that are not saturated. On the other hand, saturated pixels satisfy the constraint given by

$$(\mathbf{y}_i)^S \geq (\mathbf{k}_i * \mathbf{x})^S, \quad i = 1, \ldots, Q,$$

where $(.)^S$ represents the set of all pixels that are saturated. Our goal is to estimate the latent image $\mathbf{x}$ and the blur kernels $\{\mathbf{k}_1, \ldots, \mathbf{k}_Q\}$ under a Gaussian noise model on the measurement noise ω_i. All blurred images are assumed to be of the same size $M \times M$, pixels. The size of the blur kernels increase proportional to the exposure duration. In particular, the short-exposure images will exhibit very little blur, while the longer-exposure images will exhibit larger blurs. We assume that the blur kernel in the largest exposure image is of size $K \times K$ and for the shorter exposure images they are proportionally smaller. Further, as we move from images of short exposure duration to those of longer exposure duration, more and more pixels fall into the set of saturated pixels thereby resulting in an increasing number of inequality constraints.

Overview of the solution: As before, we employ a sparse image-gradient prior for the latent image $\mathbf{x}$ (realized using a minimum total-variation (TV) norm) and a sparse prior on the kernels $\mathbf{k}_i$ (realized using a minimum ℓ_1 norm). Under these priors, the HDR deblurring problem can be reduced to solving the following opti-

(a) Blurry image (b) Cho and Lee 2009 (c) Shan et al. 2008 (d) BlurBurst

Fig. 14. Telephoto imaging of an animal with a 300mm lens. Camera shake during telephoto imaging can lead to large amounts of blur. Here, we obtained 5 images and processed them at 512×512 resolution. (a) One of the 5 blurry input images. (b, c) Deblurred results using single-image deblurring algorithms of Cho and Lee [2009] and Shan et al. [2008]. Shown is the best of the 5 deblurred images. (d) Result obtained using BlurBurst. In each column, the top row corresponds to the image and shown below are select patches.

(a) Cho and Lee 2009 (b) Shan et al. 2008 (c) Sroubek Milanfar 2012 (d) BlurBurst

Fig. 15. Low-light imaging with a exposure of 1.6 second. The long exposure required to capture low-light scenes often leads to large blur kernels. Here, we obtained 6 images and processed them at 512×512 resolution. (a-b) Deblurring results of state-of-the-art single image deblurring algorithms. In each case, shown are the best of the 6 deblurred images. (c) Deblurring results of a state-of-the-art multiple image deblurring algorithm. (d) Result obtained using BlurBurst. In each column, the top row corresponds to the image and shown below are select patches.

mization problem:

$$\begin{aligned} &\arg\min_{\mathbf{x},\mathbf{k}_i} \mathrm{TV}(\mathbf{x}) + \lambda \sum_i \|\mathbf{k}_i\|_1 \\ &\text{subject to} \quad (\mathbf{k}_i * \mathbf{x})^S \geq 1 \\ \text{and} \quad &\sum_i \|A(\mathbf{y}_i^{NS} - (\mathbf{k}_i * \mathbf{x})^{NS})\|_2 \leq \epsilon \end{aligned} \tag{4}$$

where $A(\cdot)$ is the same high-order noise model that we discussed earlier. The constrained optimization problem, as before, in (4) is bi-convex. As before, we solve it as a sequence of convex problems — at each step, optimizing over $\mathbf{x}$ or $\mathbf{k}_i$ with the other variables held fixed at the latest estimates. The main difference is that at each step the corresponding unconstrained optimization problem is turned into a constrained optimization problem with the inequality constraints in the saturation equations providing the constraints.

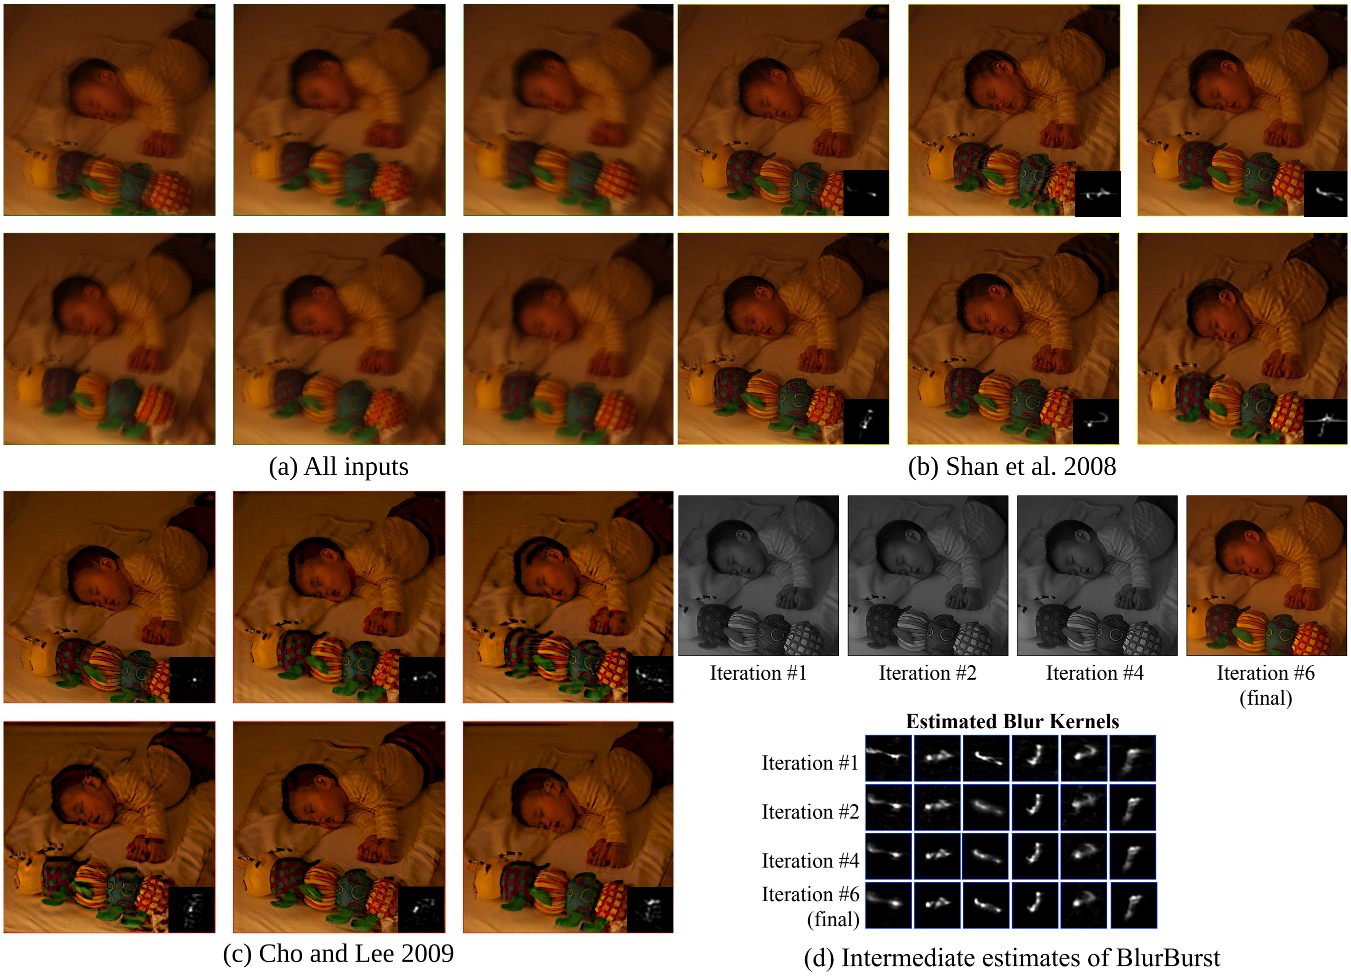

Fig. 16. Low-light imaging result. Shown are (top-left) input blurry images, results using single-image deblurring algorithms of (top-right) Shan et al. and (bottom-left) Cho and Lee. Finally, (bottom-right) shown are intermediate estimates of latent image and blur kernels from BlurBurst.

Fig. 18. Deblurring in the presence of saturation. (Left) A blurred image with saturated regions. (Right) Deblurred image using the non-blind deblurring algorithm of Shan et al. [2008]. While regions far-away from saturated regions are well-resolved (blur inset), we observe observe severe ringing artifacts around saturated regions (red inset).

We follow the same initialization and registration strategies outlined in Section 3. There are subtle differences due to the individual images being of different exposures which require us to ignore saturated regions during image registration. Another subtle difference is that the blur kernel sizes are chosen differently across images of different exposure durations.

Simulations: We show multiple examples of HDR imaging using simulations of hand-help imaging scenarios. Figure 19 showcases results in a non-blind imaging scenario, and Figure 20 showcases result in a blind imaging scenario. In both cases, we simulated 7 images, each with a different exposure time with a PSF whose size is directly proportional to the exposure duration. The blur kernel of the largest exposure was 15×15 pixels. The point spread functions were simulated using the procedure described before in Section 4. The exposure time of each input image was twice the exposure time of the previous image (e.g. 1/480, 1/240, 1/120, 1/60, 1/30, 1/15, and 1/7.5 second). Random noises of standard deviation σ =3/255 were added to the captured images. We used BlurBurst — suitably modified to handle saturation — to reconstruct the HDR image from the multiple exposure stack. As seen in the two Figures, the modified BlurBurst algorithm is successful in both cases without exhibiting any of the artifacts in competing methods.

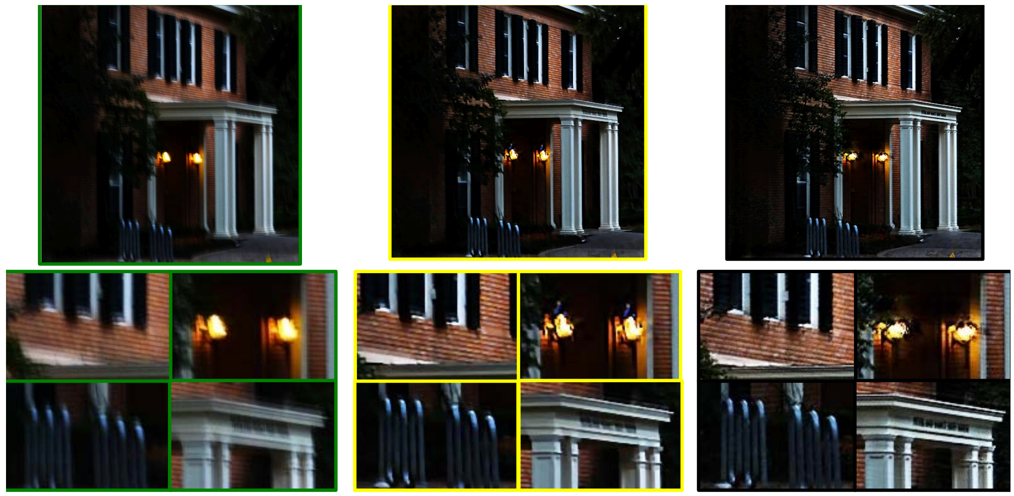

(a) Blurry image (b) Shan et al. 2009 (c) BlurBurst

Fig. 17. Low-light imaging with a exposure of 0.25 second. The long exposure required to capture low-light scenes often leads to large blur kernels. Here, we obtained and processed 8 images at 512×512 resolution. (a) One of the 8 blurry input images. (b-c) Deblurred results using single-image deblurring algorithm of Shan et al. [2008]. In each case, shown are the best of the 8 deblurred images. (d) Result obtained using BlurBurst. In each column, the top row corresponds to the image and shown below are select patches.

Real dataset: Shown in Figure 21 is a result of reconstructing the HDR image from images collected using a hand-help Canon SLR. We collected multi exposed 9 images using a Canon camera operating at ISO 100, F/6.3 and focal length 100mm. The exposure times of each image were $1/250$, $1/120$, $1/60$, $1/30$, $1/15$, $1/8$, $1/4$, $1/2$, 1 second. Their resolution in input and processing were 370×280. It goes without saying that this is blind situation and the size of blur kernels estimated was 25×25 pixels. As before, we compare our results with other HDR compositions including the input images without deblurring, and the input images individually deblurred using the algorithm in Shan et al. The HDR image recovered using BlurBurst restores fine textures with fewer artifacts in spite of many saturated areas in long exposed images.

7. DISCUSSIONS

In this paper, we have proposed a new multi-image deblurring algorithm called BlurBurst to overcome the limitation of single-image deblurring, especially in the context of large blurs due to camera shake. Obtaining a few blurry images opportunistically provides blur profiles that are not aligned; this makes the deblurring problem well-conditioned. In addition to a suite of simulation results, we have demonstrated the benefits of multi-image deblurring for two real world applications: tele-photo and low-light imaging. Finally, we have also extended our algorithm to handle deblurring in the presence of saturation to obtain hand-held HDR images.

Limitations: We discuss three key limitations of the BlurBurst framework. First, the assumptions of spatially invariant blur do not hold in all cases, including moving objects in the scene, defocus blur due to large apertures, and complex camera motion. Second, the assumption of a single latent image is violated when we have moving objects; here, background regions are selectively occluded and revealed in different images. Third, the success of our algorithm is dependent on the initial registration of the blurred images. The homography model assumed in our framework is violated for complex camera shake coupled with scenes with complex geometry.

Future work: A key finding of this paper is that multi-image *non-blind* deblurring is well-conditioned *even* in the absence of image priors. A theoretical treatment under various image priors and in the blind setting is an extremely interesting and fruitful direction for future research. Finally, inspired from ideas in [Farsiu et al. 2004], an exciting avenue for future research is towards enabling super-resolution of the latent image during the process of deblurring. We look forward to this future work.

ACKNOWLEDGMENTS

This research was supported in part by Sony Corporation. ACS was supported by the NSF grant CCF-1117939. AV was supported by NSF grants IIS-1116718 and CCF-1117939, and by a research collaboration grant from Sony Corporation. RGB was supported by the grants NSF CCF-0431150, CCF-0728867, CCF-0926127, CCF-1117939, ARO MURI W911NF-09- 1-0383, W911NF-07-1-0185, DARPA N66001-11-1-4090, N66001-11-C-4092, N66001-08-1-2065, ONR N00014-12-1-0124 and AFOSR FA9550-09-1-0432.

REFERENCES

AGRAWAL, A., XU, Y., AND RASKAR, R. 2009. Invertible motion blur in video. *ACM Trans. Graph. 28,* 3, 1–8.

BECK, A. AND TEBOULLE, M. 2009. Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems. *IEEE Trans. Image Proc. 18,* 11, 2419–2434.

BEN-EZRA, M. AND NAYAR, S. K. 2003. Motion deblurring using hybrid imaging. In *CVPR*.

CHATTERJEE, P., JOSHI, N., KANG, S. B., AND MATSUSHITA, Y. 2011. Noise suppression in low-light images through joint denoising and demosaicing. In *CVPR*.

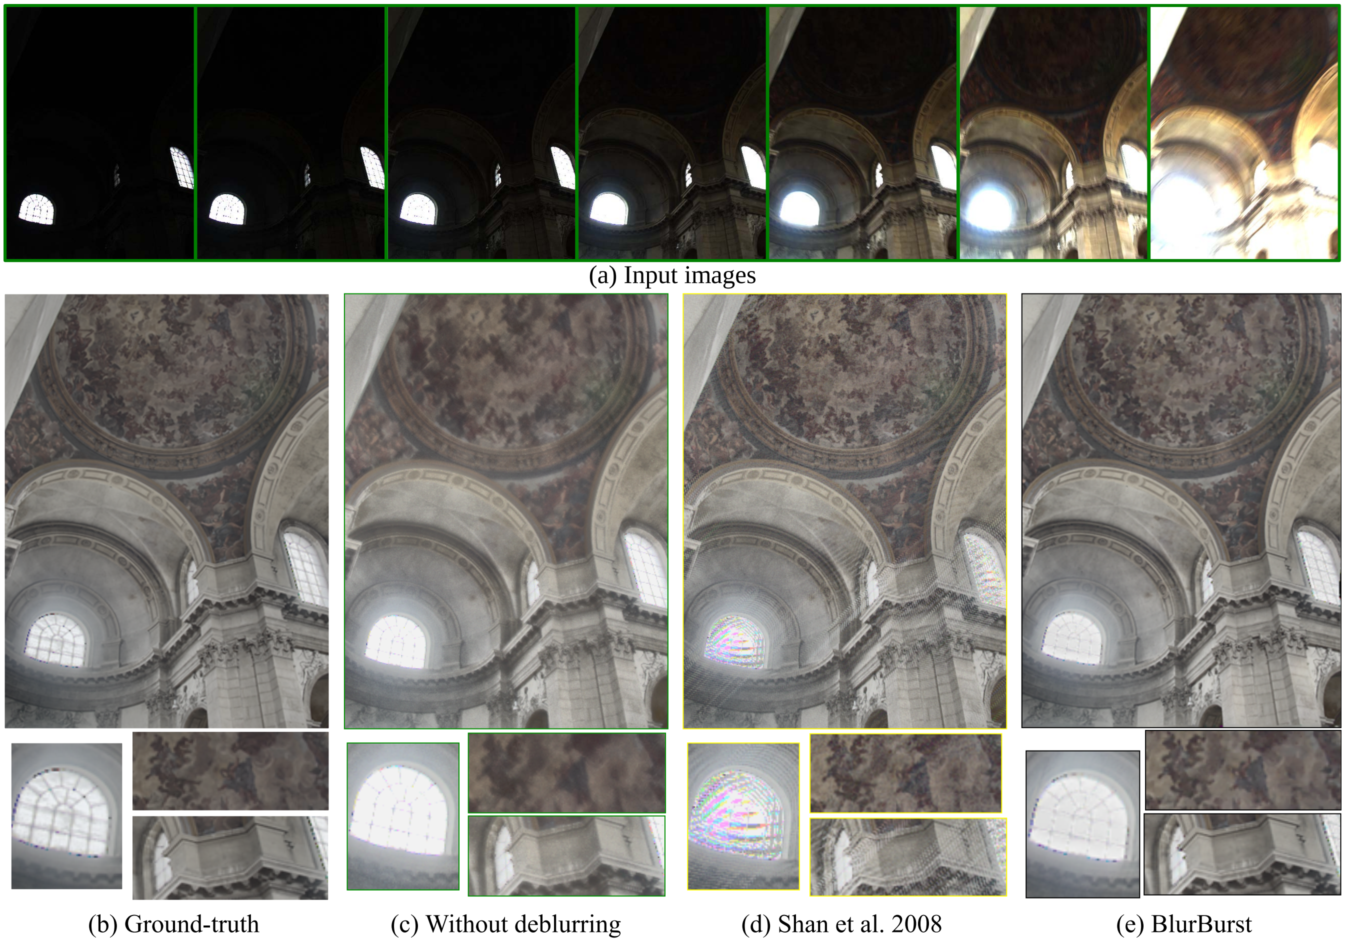

Fig. 19. Simulation of hand-held HDR imaging. (a) Input images taken with different exposure durations. We all processed at 384×512 resolution and blurred with 15×15 kernel size at most corresponding to simulated exposure durations. In this experiment, we assumed knowledge of the blur kernel. Shown are (b) HDR ground-truth data, (c) HDR composite image without deblurring, (d) HDR composite image reconstructed from inputs individually deblurred by non-blind deconvolution of Shan et al. [2008], and (e) result synthesized by non-blind version of BlurBurst.

CHO, S. AND LEE, S. 2009. Fast motion deblurring. *ACM Trans. Graphics 28,* 5, 1–8.

CHO, S., WANG, J., AND LEE, S. 2011. Handling outliers in non-blind image deconvolution. In *ICCV*.

CHO, T. S., PARIS, S., HORN, B. K. P., AND FREEMAN, W. T. 2011. Blur kernel estimation using the radon transform. In *CVPR*.

DEBEVEC, P. E. AND MALIK, J. 1997. Recovering high dynamic range radiance maps from photographs. In *SIGGRAPH*. 369–378.

FARSIU, S., ROBINSON, M., ELAD, M., AND MILANFAR, P. 2004. Fast and robust multiframe super resolution. *IEEE Trans. Image Proc. 13,* 10, 1327–1344.

FERGUS, R., SINGH, B., HERTZMANN, A., ROWEIS, S. T., AND FREEMAN, W. 2006. Removing camera shake from a single photograph. *ACM Trans. Graphics 25*, 787–794.

FISCHLER, M. A. AND BOLLES, R. C. 1981. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM 24,* 6, 381–395.

HARMELING, S., SRA, S., HIRSCH, M., AND SCHLKOPF, B. 2010. Multiframe blind deconvolution, super-resolution, and saturation correction via incremental em. In *ICIP*.

JOSHI, N. AND COHEN, M. F. 2010. Seeing Mt. Rainier: Lucky imaging for multi-image denoising, sharpening, and haze removal. In *ICCP*.

JOSHI, N., KANG, S. B., ZITNICK, C. L., AND SZELISKI, R. 2010. Image deblurring using inertial measurement sensors. *ACM Trans. Graphics 29,* 4, 1–9.

LEVIN, A., FERGUS, R., DURAND, F., AND FREEMAN, W. 2007. Image and depth from a conventional camera with a coded aperture. *ACM Trans. Graphics 26,* 3.

LEVIN, A., SAND, P., CHO, T. S., DURAND, F., AND FREEMAN, W. T. 2008. Motion-invariant photography. *ACM Trans. Graph. 27,* 3.

LIU, X. Q. AND ABBAS, E. G. 2003. Synthesis of high dynamic range motion blur free image from multiple captures. *IEEE Trans. Circuits and Systems 50,* 4, 530–539.

LU, P. Y., HUANG, T. H., WU, M. S., CHENG, Y. T., AND CHUANG, Y. Y. 2009. High dynamic range image reconstruction from hand-held cameras. In *CVPR09*.

LUCY, L. B. 1974. Iterative technique for rectification of observed distributions. *Astronomical Journal 79,* 6, 745–754.

Fig. 20. Simulation of hand-held HDR imaging. (a) Input images taken with different exposure durations. We all processed at 512×384 resolution and blurred with 15×15 kernel size at most corresponding to simulated exposure durations. In this experiment, we assumed that each blur kernels was unknown and we iteratively estimated them to synthesize our result. Shown are (b) HDR ground-truth data, (c) HDR composite image without deblurring, (d) HDR composite image reconstructed from inputs individually deblurred by blind deconvolution of Shan et al. [2008], and (e) result synthesized by BlurBurst.

NEELAMANI, R., CHOI, H., AND BARANIUK, R. G. 2004. ForWaRD: Fourier-wavelet regularized deconvolution for ill-conditioned systems. *IEEE Trans. Signal Proc. 52,* 2, 418–433.

OSHER, S., BURGER, M., GOLDFARB, D., XU, J., AND YIN, W. 2005. An iterative regularization method for total variation-based image restoration. *Multiscale Modeling and Simulation 4,* 2, 460–489.

POOR, H. 1988. *An introduction to signal detection and estimation*. Springer-Verlag.

RASKAR, R., AGRAWAL, A., AND TUMBLIN, J. 2006. Coded exposure photography: Motion deblurring using fluttered shutter. In *SIGGRAPH*.

RAV-ACHA, A. AND PELEG, S. 2005. Two motion-blurred images are better than one. *Pattern Recogn. Lett. 26,* 3, 311–317.

RICHARDSON, W. H. 1972. Bayesian-based iterative method of image restoration. *J. Optical Society of America 62,* 1, 55–59.

SHAN, Q., JIA, J., AND AGARWALA, A. 2008. High-quality motion deblurring from a single image. *ACM Trans. Graphics 27*, 1–10.

SHEIKH, H. R. AND BOVIK, A. C. 2006. Image information and visual quality. *IEEE Trans. Image Proc. 15,* 2, 430–444.

SROUBEK, F. AND FLUSSER, J. 2003. Multichannel blind iterative image restoration. *IEEE Trans. Image Proc. 12,* 9, 1094–1106.

SROUBEK, F. AND FLUSSER, J. 2005. Multichannel blind deconvolution of spatially misaligned images. *IEEE Trans. Image Proc. 14,* 7, 874–883.

SROUBEK, F. AND MILANFAR, P. 2012. Robust multichannel blind deconvolution via fast alternating minimization. *IEEE Trans. Image Proc. 21,* 4, 1687–1700.

VAN DEN BERG, E. AND FRIEDLANDER, M. P. 2008. Probing the pareto frontier for basis pursuit solutions. *SIAM J. Scientific Computing 31,* 2, 890–912.

WANG, Z., BOVIK, A. C., SHEIKH, H. R., AND SIMONCELLI, E. P. 2004. Image quality assessment: From error visibility to structural similarity. *IEEE Trans. Image Proc. 13,* 4, 600–612.

YUAN, L., SUN, J., QUAN, L., AND SHUM, H.-Y. 2007. Image deblurring with blurred/noisy image pairs. *ACM Trans. Graph. 26,* 3.

YUAN, L., SUN, J., QUAN, L., AND SHUM, H.-Y. 2008. Progressive inter-scale and intra-scale non-blind image deconvolution. *ACM Trans. Graph. 27,* 3.

(a) Input Images

(b) Without deblurring

(c) Shan et al. 2008

(d) BlurBurst

Fig. 21. HDR image estimated by deblurring multiple images acquired from a hand-held SLR. (a) Four input images, from a subset of the 9 images. taken with different exposure durations. We obtained and processed them at 370×280 resolution. (b-d) We estimate three different HDR composite images using different techniques for deblurring. The HDR composite recovered using BlurBurst has very little artifacts handling as compared to composites created without deblurring and using single-image deblurring method of Shan et al. [2008].