

Resolving Objects at Higher Resolution from a Single Motion-blurred Image

Amit Agrawal, Ramesh Raskar

TR2007-036 July 2007

Abstract

Motion blur can degrade the quality of images and is considered a nuisance for computer vision problems. In this paper, we show that motion blur can in-fact be used for increasing the resolution of a moving object. Our approach utilizes the information in a single motion-blurred image without any image priors or training images. As the blur size increases, the resolution of the moving object can be enhanced by a larger factor, albeit with a corresponding increase in reconstruction noise. Traditionally, motion deblurring and super-resolution have been ill-posed problems. Using a coded-exposure camera that preserves high spatial frequencies in the blurred image, we present a linear algorithm for the combined problem of deblurring and resolution enhancement and analyze the invertibility of the resulting linear system. We also show a method to selectively enhance the resolution of a narrow region of high-frequency features, when the resolution of the entire moving object cannot be increased due to small motion blur. Results on real images showing up to four times resolution enhancement are presented.

IEEE CVPR 2007

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Resolving Objects at Higher Resolution from a Single Motion-blurred Image

Amit Agrawal and Ramesh Raskar
Mitsubishi Electric Research Labs (MERL)
201 Broadway, Cambridge, MA, USA 02139
[agrawal, raskar]@merl.com

Abstract

Motion blur can degrade the quality of images and is considered a nuisance for computer vision problems. In this paper, we show that motion blur can in-fact be used for increasing the resolution of a moving object. Our approach utilizes the information in a single motion-blurred image without any image priors or training images. As the blur size increases, the resolution of the moving object can be enhanced by a larger factor, albeit with a corresponding increase in reconstruction noise.

Traditionally, motion deblurring and super-resolution have been ill-posed problems. Using a coded-exposure camera that preserves high spatial frequencies in the blurred image, we present a linear algorithm for the combined problem of deblurring and resolution enhancement and analyze the invertibility of the resulting linear system. We also show a method to selectively enhance the resolution of a narrow region of high-frequency features, when the resolution of the entire moving object cannot be increased due to small motion blur. Results on real images showing up to four times resolution enhancement are presented.

1. Introduction

A limited resolution sensor may fail to recover sufficient details on a static object. If we let the object move causing motion blur, the resulting smear distributes the signal information to neighboring pixels. Can we exploit these additional pixel samples to reconstruct the object at a higher resolution? Unfortunately, both motion deblurring and resolution enhancement from multiple images are ill-posed problems [1, 22, 4] for photos captured with a traditional camera.

Can motion-blur be used to enhance the resolution of a moving object? In this paper, we show that if the blurred image is larger than the size of the moving object in pixels by a factor s , the resolution of the moving object can also be enhanced upto a factor s (Figure 1). To practically achieve this, we use a modified coded-exposure camera that preserves high spatial frequencies in the motion blurred im-




Figure 1. Can we achieve resolution enhancement by letting objects blur? (a) Camera pixel resolution is insufficient to capture individual alternating lines in a *static* image of a standard resolution chart since they span only $n = 18$ pixels. (b) Zoom showing four times upsampling in horizontal direction. (c) If we intentionally let the object blur horizontally, the black lines occupy more pixels. (d) The crop shows that now the smeared features spans $t = 74$ pixels with motion blur $k = 57$ pixels. (e) The additional samples support deblurring with a resolution enhancement factor of $s = \lfloor \frac{74}{18} \rfloor = 4$. (f) and (g) For comparison, results of traditional deblurring and its four times upsampling fails to recover the details. (h) Zoomed in close up of the static resolution chart.

age.

Contributions: The contributions of this paper are as follows.

- We show that the resolution of a moving object can be enhanced by utilizing information from motion-blur, if the blur size is sufficiently large.
- In practice, resolution enhancement is ill-posed and produces meaningless reconstruction in the presence of noise. We show that by using a coded-exposure camera, one can improve the invertibility and achieve practical resolution enhancements.
- We analyze the relationship between the blur size and parameters of the code and show how to choose these parameters for a given blur size and resolution enhancement factor (REF).
- We propose a method to selectively increase the resolution of a high frequency region on an object when the blur is small.

Our procedures and **assumptions** are as follows. We capture a *single* photo of a dynamic scene from a static camera. We assume lens optical blur is smaller than one pixel (in the high resolution space) and also assume that the moving object is in the plane of focus. For simplicity, we restrict our discussion to 1-D linear motions of the object, although the technique can be extended to more complex object motions and 2D motions such as camera shake. Thus, the resolution can be enhanced only along the direction of the motion. We assume either a known point spread function (PSF) for the motion blur or that it can be manually specified by describing the direction and extent of motion for 1-D linear motions. For practical deblurring, we assume that the appearance of the object during the exposure time remains invariant. Thus, specular, transparent and other view-dependent appearance changes can distort the results.

1.1. Related work

Super resolution (SR): SR refer to techniques aimed at improving the resolution of a low-resolution sensor. Numerous algorithms which combine multiple low-resolution images into a single high resolution image have been proposed in literature [24, 7, 13, 19, 10, 16, 9, 11, 25, 2]. In these approaches, first the relative motion between the camera and the scene is estimated and all images are registered to a reference frame. Then the images are fused to obtain a high resolution image. In [4], it was argued that motion blur in individual images significantly degrades the quality of super-resolution. They propose a jitter camera, where the sensor was modified to obtain four images using controlled sub-pixel detector shifts. A text-specific bimodal prior for super-resolution of text in videos was proposed in [6].

High resolution recovery of 1-D features was proposed in [5]. However, their method is restricted as they assume

that each row or column of the low resolution image results from the sampling of the same signal with different shifts (*e.g.* bar codes). In contrast, our approach has no such restriction and the profiles of 1-D signal on the object can be different along the direction orthogonal to the motion direction.

Single image SR is an under-constrained problem and most of the previous work can be divided into (a) reconstruction based algorithms, where a high resolution image is sought so that after downsampling, it is as close as possible to the low resolution input image, (b) learning based Bayesian methods [21, 23] using training dataset and image priors and (c) functional interpolation which results in blurring of the discontinuities. Our approach is different from algorithms in all the above three categories and is not a hallucination algorithm. We show the resolution enhancement using motion blur is not an under-constrained problem if the blur is large and present a linear algorithm which does not use any image priors or training data.

Motion deblurring: Image deblurring or deconvolution [12] has been a well-studied problem. See [14] for a survey in this area. Blind image deconvolution [15] attempts to infer the sharp image and the PSF simultaneously from the given image, based on various assumptions applied to PSF. Bayesian methods assume specific image priors such as the Poisson distribution as in Richardson-Lucy algorithm [18]. Ben-Ezra & Nayar [3] proposed a hybrid camera where a low resolution video camera was used to estimate the PSF, which was then used to deblur high resolution image from a digital still camera.

Coded sampling: Methods to preserve spatial frequencies for subsequent reconstruction include coded aperture imaging [20] used in astronomy to overcome the limitations of a pinhole camera. Modified Uniformly Redundant Arrays (MURA) [8] are used for coding and decoding the light distribution of distant stars using circular convolution and deconvolution. A coded exposure camera [17] can preserve high spatial frequencies in a motion-blurred image and make the deblurring process well-posed.

2. Resolution enhancement using motion blur

We first formulate the problem of resolving a moving object at higher resolution using motion blur. We represent quantities at high resolution using subscript s . Let u_s denote the image of the object captured using a virtual high resolution sensor if the object was static. The observed low resolution motion blurred image \mathbf{b} is modeled as

$$\begin{aligned} \mathbf{b} &= h_{sensor} * h_{motion} * h_{lens} * \mathbf{u}_s + \boldsymbol{\eta}, \\ &= h_{sensor} * h_{motion} * \mathbf{x}_s + \boldsymbol{\eta}, \end{aligned} \quad (1)$$

where $*$ denotes convolution, $\boldsymbol{\eta}$ denotes the noise in the imaging process and h_{lens} , h_{motion} and h_{sensor} denote the lens, motion and sensor PSF's respectively. We assume




Figure 2. Basic concept of enhancing resolution using motion blur. (a) (Top) In traditional deblurring, the motion blurred image obtained from sensor is deblurred using h_{motion} at the current resolution. It may fail to recover high spatial frequencies on the object if the resolution is low. (Bottom) Our approach inverts the combined effect of h_{motion} and h_{sensor} by modeling the blurring process at higher resolution followed by decimation. This requires blur to be large enough so that $t > n_s$. In practice, this linear system is ill-posed. (b) By using a coded exposure camera, the smear matrix A_s changes to a banded matrix according to the binary code used for fluttering the shutter. This makes (2) well-posed and invertible. (c) Estimated \mathbf{x}_s by adding noise in \mathbf{b} and solving (2) using least squares (matlab backlash operator). For the coded exposure camera, the condition number of the K_s matrix is orders of magnitude lower, facilitating faithful recovery of \mathbf{x}_s .

that CCD pixels are square as in [1] and thus h_{sensor} is a square PSF related to the pixel size. h_{motion} depends on object motion within the exposure time of the captured image. Specifically, by improving resolution in this paper, we refer to inverting the combined effects of h_{sensor} and h_{motion} . We assume lens PSF is less than one pixel in the high resolution grid, or resolution can be improved only upto the lens PSF. Thus, $\mathbf{x}_s = (h_{lens} * \mathbf{u}_s)$ is our desired high resolution static image of the object. In addition, we assume that the object moves in a straight line with constant speed within the exposure time, leading to a spatially invariant PSF h_{motion} .

Resolution enhanced deblurring: Let s denote the resolution enhancement factor (REF) and k denote the size of blur (in pixels) in the low resolution sensor output. Without loss of generality, assume that the object is moving vertically. Let n and $n_s = n * s$ denote the height of the object in low resolution and high resolution grid respectively and w denotes its width. Then each column of the observed image \mathbf{b} will be smeared in the vertical direction and will have $t = k + n - 1$ samples due to linear convolution.

Discretizing (1) relates the high resolution image \mathbf{x}_s to the low resolution blurred image \mathbf{b} as

$$\begin{aligned} D_s * A_s * \mathbf{x}_s &= \mathbf{b}, \\ K_s * \mathbf{x}_s &= \mathbf{b}, \end{aligned} \quad (2)$$

where \mathbf{x}_s is $n_s \times w$ high-resolution object, A_s is the smearing matrix representing motion blur in the high resolution grid and D_s is the decimation matrix due to h_{sensor} . Each row of D_s has only s ones, with each column having at most a single one. Each column of A_s corresponds to a circularly shifted PSF vector of length $s * k$ zero-padded to the size of \mathbf{x}_s , representing linear convolution¹. $K_s = D_s * A_s$ denote the combined deblurring and decimation matrix. Figure 2(a)

¹Note that the motion blur in high resolution grid will be $s \times k$

shows a toy example for $s = 2$, where a 16 pixel tall high frequency pattern moves vertically for $k = 50$ pixels.

Traditional deblurring: In traditional deblurring, one attempts to recover \mathbf{x} at current resolution by solving

$$A\mathbf{x} = \mathbf{b}, \quad (3)$$

where A denotes the smear matrix for blur size k . Thus, traditional deblurring only tries to invert the effect of h_{motion} at the current resolution. On the other hand, our approach inverts the combined effect of h_{motion} and h_{sensor} by modeling the motion blur process in high resolution, followed by decimation to account for low resolution. In fact, (2) can also be written as

$$D_s * A_s * \mathbf{x}_s = D_s * \mathbf{b}_s = \mathbf{b}, \quad (4)$$

where \mathbf{b}_s is the motion-blurred image in high resolution.

A simple counting argument shows that the number of observations t in the low resolution motion blurred image should be greater than the number of unknowns n_s for resolution enhancement. Thus

$$\begin{aligned} t &\geq n_s \\ (n + k - 1) &\geq n * s \\ k &\geq n * (s - 1) + 1 \end{aligned} \quad (5)$$

Thus, motion blur k has to be larger than the size of the object in the direction of the motion by a factor of $s - 1$ for resolution enhancement by a factor of s . For example, if $s = 4$, the object should move more than three times its size in the direction of the motion. In Figure 1, the chart with width $n = 18$ pixels blurs by $k = 57$ pixels, more than three times its width. This facilitates resolution enhancement by a factor of four. Thus, as the size of the blur k increases, REF can be increased. However, in practice, this results in more noise in the estimated image as shown by the analysis in Section 3.2.

3. Invertibility of resolution enhancement

In practice, the linear system (2) is ill-posed. Previous research have shown that deblurring and super-resolution are ill-posed problems. Baker & Kanade [1] analyzed the theoretical limits on super-resolution from multiple images, assuming that the images are perfectly registered. They showed that for square PSF, the reconstruction constraints are not invertible, and the condition number for integer magnification factors (> 1) is infinity. The linear system is not invertible because high frequencies are lost in the sensor output due to motion blur and low resolution, resulting in zeros in the frequency spectrum of the PSF's. Tanaka & Okutomi [22] presented a condition number theorem for super-resolution and analyzed Gaussian and box PSF. Ben-Ezra *et al.* [4] showed that any motion blur in individual images is bad for super-resolution using multiple images.

3.1. Coded exposure imaging

One way to preserve high spatial frequencies in a motion-blurred image is to use a coded exposure camera [17]. We analyze the invertibility of the linear system for resolution enhancement and show that it is significantly improved using a coded exposure camera.

The coded exposure camera is a simple modification of a conventional camera in terms of the integration of light within a single frame (see Figure 3). Instead of keeping the shutter open for the entire exposure duration, a coded exposure camera "flutters" the shutter open and closed according to a carefully chosen pseudo-random binary sequence within a single frame. For static parts of the scene, this only results in reducing the effective exposure time by half. However, for moving objects, the resulting motion-blurred image exhibits *coded* blur as opposed to continuous blur for a traditional camera. The blurred image is a convolution of the static image with the temporal code used for modulating the exposure. Effectively, the motion PSF is changed from a low pass box filter to a broad-band filter. The frequency response of this broadband filter depends on the binary code used for fluttering the shutter. Since the resulting motion PSF becomes broad-band and does not have zeros in its frequency spectrum, resolution enhancement becomes a well-posed invertible problem.

In terms of the linear system (2), the A matrix is changed if the exposure is coded as shown in Figure 2. For a traditional camera, the motion PSF is a box filter, represented as all ones² in each column of A . By coding the exposure, this is changed to the binary code used for fluttering.

How does coding the exposure help in solving the system? Let us compare the estimated \mathbf{x}_s in presence of noise in \mathbf{b} . We add zero mean Gaussian noise ($\sigma = 0.1$) to the blurred images \mathbf{b} shown in Figure 2. Figure 2(c) shows the estimated \mathbf{x}_s using traditional and coded exposure camera

²assuming that the object is moving at a constant speed




Figure 3. Toy example showing the difference between a traditional and coded exposure camera. In a traditional camera, the shutter is kept open for the entire exposure time resulting in loss of high spatial frequencies and a continuous blur in the blurred image. In a coded exposure camera, the shutter is *fluttered* (open and closed) according to a binary code within the exposure time, preserving high spatial frequencies in the blurred image.

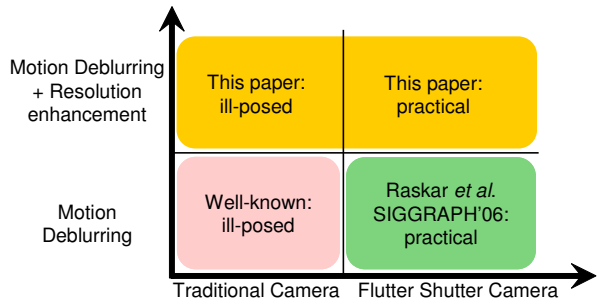


Figure 4. Motion deblurring using a traditional camera is a well-known but ill-posed problem. The flutter shutter camera [17] makes deblurring a well-posed problem. Our work differs in that we combine deblurring with resolution enhancement. We show that in practice, the resulting linear system remains ill-posed using a traditional camera. By using a flutter shutter camera, the conditioning of the system can be improved.

by solving the linear system in (2). Note that the condition number of the K_s matrix for traditional camera is of the order of 10^{18} . In contrast, the condition number of K_s using coded exposure camera is only 149.2. Thus, even with small amount of noise, (2) is non-invertible for a traditional camera but becomes invertible using a coded exposure camera. Figure 4 summarizes the differences with this work, traditional motion deblurring and coded exposure based deblurring [17].

3.2. Theoretical analysis

Consider a length m binary code according to which the camera's shutter is open and closed within one exposure time. We study the relationship between the code length m , blur size k , and REF, s . We analyze (2) in terms of the noise covariance of the estimated \mathbf{x}_s . For a linear system $\mathbf{Ax} = \mathbf{b}$, assuming IID Gaussian noise in \mathbf{b} with mean 0 and variance σ^2 , the covariance matrix of the estimated $\hat{\mathbf{x}}$ is given by $\sigma^2 \Sigma$, where $\Sigma = (A^T A)^{-1}$. The mean square error (MSE) of $\hat{\mathbf{x}}$ is $\frac{\sigma^2}{n} \text{Trace}(\Sigma)$. Thus, the signal to noise ratio (SNR) can be defined as $\frac{n}{\text{Trace}(\Sigma)}$.

Blur size and code length for $s = 1$: We first analyze




Figure 5. Relationship between code length m , blur size k and REF, s . (a) Plots of SNR versus relative blur size k/m for coded exposure camera using codes of different lengths for traditional deblurring ($s = 1$). Note that for each curve, SNR degrades after $k > m$. (b) SNR for a coded exposure camera is much better compared to a traditional camera for wide range of code lengths and motion blur sizes. (c) Similar plots as in (b) for combined deblurring and resolution enhancement ($s = 2$), showing that it is impractical using a traditional camera. (d) SNR versus s for a fixed blur size k ($= 100$) pixels and object size n ($= 25$) pixels. As s increases, reconstruction noise also increases.

the relationship between blur size k and code length m for traditional deblurring ($s = 1$). Figure 5(a) shows plots of SNR (in db) versus *relative* blur size k/m for fixed object size ($n = 25$ pixels) for coded exposure camera using different code lengths m . For each m , a good code was found as described in Section 3.3. Note the for each curve, SNR remains high till k equals m ($\log(k/m) = 0$). Thus, a code of length m is the most effective when the blur size k is equal to m . Figure 5(b) shows same plots as in Figure 5(a) along with the SNR plot for a traditional camera. The X axis is changed to blur size k for fair comparison with the traditional camera. For small blur ($k \approx 10$ pixels), coded-exposure camera performs similar to a traditional camera. As the blur size increases, the performance of coded-exposure camera improves significantly. Note that for $k \leq 52$, SNR is highest for $m = 52$ (red curve), for $50 < k \leq 100$, SNR is highest for $m = 100$ (green curve), emphasizing that code length m should be close to blur size k .

Blur size and code length for $s = 2$: Figure 5(c) shows SNR plots similar to Figure 5(b) but for $s = 2$. The SNR for a traditional camera degrades sharply even for small motion blur, making it impractical for resolution enhancement. In comparison, the coded-exposure camera makes the linear system (2) more stable. For code length m , the corresponding curve peaks at $k \approx \frac{m}{s}$. This is because the effective blur in high resolution space is $k * s$ and the best code should be of the same length.

SNR and REF: Figure 5(d) shows the plot of SNR ver-

sus resolution enhancement factor s for fixed $k = 100$ pixels, $n = 25$ pixels and code length $m = 400$. Again, the performance of a traditional camera degrades quickly even with small values of s . Interestingly, the SNR values are low for integer values of REF, similar to the observation made in [1]. As REF increases, SNR decreases even for coded exposure camera. Thus, higher resolution enhancement leads to more noise in the estimated image. To summarize,

- By coding the exposure, the invertibility of the linear system for deblurring and resolution enhancement is greatly improved.
- For a given enhancement factor s and blur size k , the optimal code length is $m \approx k * s$. However, mismatch in the code length does not degrade results significantly.
- Since for reasonable values of s , large motion blur is required, this emphasize the need for finding good long codes.

3.3. Fast broadband code search

The "goodness" of the code is based on maximizing the minimum of the FFT of the zero padded code so that the motion PSF does not have zeros in its frequency spectrum. We implemented a fast randomized search where a length m binary sequence is randomly generated, and stored if it is better than the previous one. Thus, one need to store only the current best sequence. This approach gives a good sequence within one hour of Matlab computation time. However, if the system specifications are known (such as the approximate blur size k and REF s), one can maximize the SNR as defined above in searching for the code.

4. Results

We use a digital Canon camera with an external ferro-electric shutter placed in front of the main lens for coding the exposure, similar to [17]. The ferro-electric shutter becomes opaque/transparent according to the binary code being 0/1 respectively. We use a 100 long code³ with total exposure time of 200 ms. We estimate the PSF (amount of blur in pixels and motion direction) manually.

We compare the resolution enhancement results with (a) deblurring result (referred to as "traditional deblurring") up-sampled by REF, s , and (b) static image upsampled by s . Bi-cubic interpolation is used for upsampling. The reader should note that the resolution is enhanced only along the direction of the motion and that all images were captured using the coded exposure camera. Thus, the "traditional deblurring" results for comparison refer to deblurring results on coded exposure images.

Figure 1 shows results on a standard resolution chart. Resolution enhancement by a factor of four was achieved.

³ 1101001001110100000100010011010000110110010010000110101011111101010100101011010000110011101100011




Figure 6. Text resolution enhancement. (a) Captured blurred photo of a bus. (b) Rectified image to make motion lines horizontal. (c) Magnified cutout corresponding to the yellow box in rectified image. (d) Estimated deblurred image \mathbf{x} at current resolution. (e) Upsampled \mathbf{x} by two in horizontal direction. (f) Estimated resolution enhanced deblurred image \mathbf{x}_s using $s = 2$. Note the separation between letters "0", "3" and "8" is clear in our result.

Our approach was able to resolve the lines upto the top of the chart as compared to the upsampled static image or upsampled deblurred result at low resolution. Figure 6 shows a challenging outdoor scenario of increasing the resolution of text on the bus. The captured image is first rectified manually to make the direction of motion horizontal. Only the cutout corresponding to the text is processed. The separation between the letters "038" is clear in the estimated \mathbf{x}_s as compared to the low resolution deblurring. Figure 7 shows a similar example of increasing the resolution of text "4010X05" on a moving cart.

5. Selective enhancement of resolution

In Section 2, we showed that for increasing the resolution by a factor of two, the object should blur equal to its size in the direction of the motion. This might be restrictive in certain scenarios. For example, one might be interested in increasing the resolution of a small patch (*e.g.* logo or text) on an object rather than the entire object. We propose to selectively enhance the resolution of an object patch to handle such cases. The key idea is that although motion blur may be small compared to the size of the object, it might be large enough compared to the size of an individual patch or region of interest on the object to facilitate resolution enhancement.

The resolution of a narrow patch P on an object can be increased selectively when the motion blur of the object is small, if it is surrounded by low frequency neighborhood regions. To achieve this, we reorganize the linear system (2)

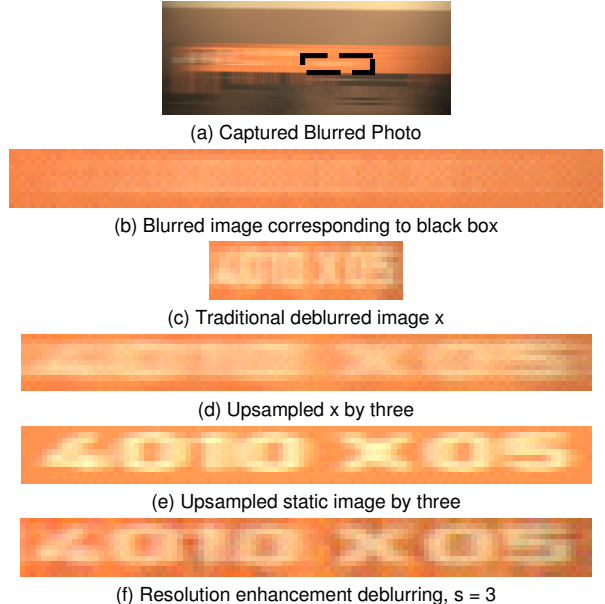


Figure 7. Text resolution enhancement. (a) Captured blurred photo. (b) Cutout corresponding to the letters "4010X05" highlighted within the black outline in (a). Motion blur $k = 77$ pixels, $t = 113$ pixels in horizontal direction. (c) Image \mathbf{x} deblurred at current resolution has $n = 37$ pixels in motion direction. (d) Upsampled image \mathbf{x} by three in the motion direction. (e) We also captured a static image of the cart. Upsampled static image by three in the motion direction. (f) Since the number of samples t is greater than $3 * n$, resolution enhancement by a factor of three can be achieved. Note that the letters "0" and "5" are clearly separated in our result.

so as to assume more unknowns on the patch P (sufficient to do resolution enhancement deblurring) and less unknowns on surrounding regions (sufficient to do low resolution deblurring). Suppose the patch P of size n^p is surrounded by patches L and R of sizes n^l and n^r respectively so that

$$n = n^l + n^p + n^r \quad (6)$$

Let $t < s * n$ in (5) so that the resolution of the entire object cannot be increased by s . In other words, there are more unknowns than observations. We first form the K_s matrix for the entire object. Then we resample the rows of K_s to reduce the number of unknowns and obtain a new matrix \hat{K}_s which is then used in (2) (See Figure 9).

We adopt a two step process. We first deblur the image to obtain \mathbf{x} . Then simple user interaction identifies patches L , P and R . If we increase the resolution of P by s , then the available samples for patches L and R are $t - s * n^p$. Dividing the number of samples equally⁴ between L and R , the available samples for each of the patches L and R are $q = (t - s * n^p) / 2$. We resample K_s where for each row (see Figure 9)

⁴Note that the samples can also be distributed according to the relative size of R and L .

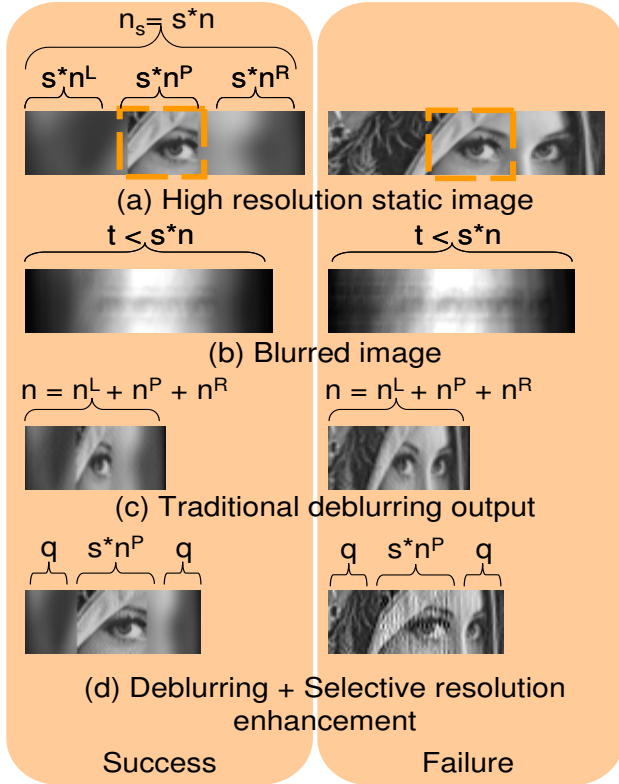


Figure 8. Selective enhancement of resolution of a narrow high-frequency patch surrounded by large low frequency neighborhoods. Synthetic example showing success (left) and failure (right) cases for $s = 2$. Image **b** simulated by blurring the high resolution image ($n_s = 128$ pixels in horizontal direction) by motion blur $k = 50$ pixels and downsampling by two has $t = 114$ pixels along motion direction. Traditional deblurring will estimate a low resolution output \mathbf{x} having $n = t - k + 1 = 65$ pixels. Since $t < 2 * n$, resolution of the entire image cannot be increased by two. However, the resolution of eye patch can be selectively enhanced by estimating two times more unknowns in the that region, while rest of the image can be deblurred at low resolution.

- The first $s * n^l$ entries corresponding to patch L are resampled to length q .
- The next $s * n^p$ entries corresponding to the patch P are kept as such.
- The last $s * n^r$ entries corresponding to patch R are resampled to length q .

Figure 8 shows a synthetic example where the static image is blurred horizontally and downsampled. Due to the small motion blur, the resolution cannot be increased by two for the entire image in the horizontal direction. However, one can selectively increase the resolution of the eye patch as shown when the neighboring regions have low frequencies. Figure 8 also shows the failure case if the surrounding patches have high frequency texture. Figure 9 shows a real example where the resolution of text "410X05" is increased by two, while rest of the image is deblurred at low resolution. Note that this example is different from the

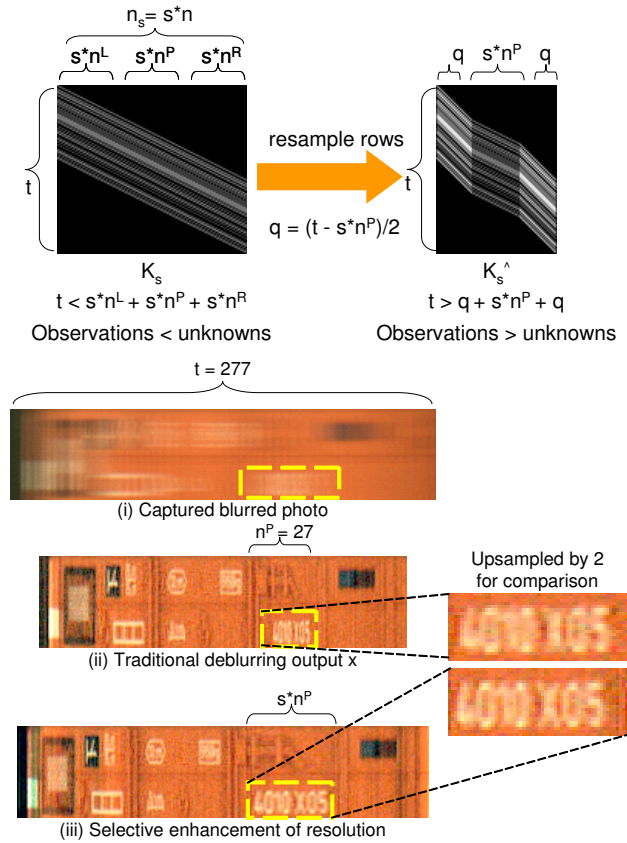


Figure 9. (Top) Resampling rows of K_s to yield \widehat{K}_s for patch of length n^p . \widehat{K}_s is then used in (2). (Bottom) Real example where the resolution of the text "410X05" is enhanced by two while the rest of the image is deblurred at low resolution. The motion blur k equals 39 pixels, which is larger than the size of the patch $n^p = 27$ pixels. Thus, REF $s = 2$ can be achieved for the patch.

one in Figure 7, where the resolution of the entire deblurred output \mathbf{x} was increased by three.

6. Discussion

Handling background: In practice, when an object moves, it smears with the background and the unknown background color along the scanline also needs to be estimated. We assume that the background is low frequency and does not have high frequency texture. We estimate a single unknown color for the background on every motion line.

Limitations and future work: A future application could be increasing the resolution of a static scene from a single low resolution image captured by intentionally shaking the camera. We have dealt with simple linear motions and thus recovery of high frequencies and resolution enhancement can only be done along the motion direction. More complicated motions can lead to spatially varying PSF. Automatic PSF estimation from a single image is a challenging problem. We currently rely on user input to in-

fer the direction of the motion and the blur size. Textured background will also create problems for deblurring. Our approach is linear and fast, although the use of specific priors such as bimodal distribution for text can improve results in those situations.

Conclusions: We have shown that motion blur can be used to enhance the resolution of a moving object from a single photo, if the blur is sufficiently large. We showed that resolution enhancement is an ill-posed problem for images captured using a traditional camera and that the coded exposure camera significantly improves the invertibility of the resulting system. Theoretical analysis of the relationship between code length, size of the blur and REF shows that our approach can handle a range of blur values and resolution enhancement factors. Results on up to four times resolution enhancement were shown to be possible, but with selective enhancement, larger enhancement factors may be achieved.

Acknowledgments

We thank anonymous reviewers and several members of MERL for their suggestions. We also thank Jack Tumblin, Joseph Katz, Keisuke Kojima, Hideaki Nii, and Haruhisa Okuda, Mitsubishi Electric, Japan, for helpful discussions and support.

References

- [1] S. Baker and T. Kanade. Limits on super-resolution and how to break them. *IEEE Trans. Pattern Anal. Machine Intell.*, 24:1167–1183, Sept. 2002. [1](#), [3](#), [4](#), [5](#)
- [2] B. Bascle, A. Blake, and A. Zisserman. Motion deblurring and super-resolution from an image sequence. In *Proc. European Conf. Computer Vision*, pages 573–582, London, UK, 1996. Springer-Verlag. [2](#)
- [3] M. Ben-Ezra and S. Nayar. Motion-based Motion Deblurring. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 26(6):689–698, Jun 2004. [2](#)
- [4] M. Ben-Ezra, A. Zomet, and S. Nayar. Video super-resolution using controlled subpixel detector shifts. *IEEE Trans. Pattern Anal. Machine Intell.*, 27:977–987, June 2005. [1](#), [2](#), [4](#)
- [5] F. Champagnat, G. Le Besnerais, and C. Kulcar. Continuous super-resolution for recovery of 1-d image features: Algorithm and performance modeling. In *Proc. Conf. Computer Vision and Pattern Recognition*, volume 1, pages 916–926, 2006. [2](#)
- [6] K. Donaldson and G. Myers. Bayesian super-resolution of text in video with a text-specific bimodal prior. In *Proc. Conf. Computer Vision and Pattern Recognition*, volume 1, pages 1188–1195, 2005. [2](#)
- [7] M. Elad and A. Feuer. Restoration of a single superresolution image from several blurred, noisy, and undersampled measured images. *IEEE Trans. Image Processing*, 6:1646–1658, Dec. 1997. [2](#)
- [8] S. R. Gottesman and E. E. Fenimore. New family of binary arrays for coded aperture imaging. *Applied Optics*, 28(20):4344–4352, Oct 1989. [2](#)
- [9] M. Irani and S. Peleg. Improving resolution by image registration. In *CVGIP*, volume 53, pages 231–239, 1991. [2](#)
- [10] M. Irani and S. Peleg. Image sequence enhancement using multiple motions analysis. In *Proc. Conf. Computer Vision and Pattern Recognition*, pages 216–221, 1992. [2](#)
- [11] M. Irani and S. Peleg. Motion analysis for image enhancement: Resolution, occlusion, and transparency. *J. of Visual Communication and Image Representation*, 4(4):324–335, Dec. 1993. [2](#)
- [12] P. Jansson. *Deconvolution of Image and Spectra*. Academic Press, 2nd edition, 1997. [2](#)
- [13] S. Kim and W.-Y. Su. Recursive high-resolution reconstruction of blurred multiframe images. *IEEE Trans. Image Processing*, 2:534–539, Oct. 1993. [2](#)
- [14] D. Kundur and D. Hatzinakos. Blind image deconvolution. *IEEE Signal Processing Magazine*, 13(3):43–64, May 1996. [2](#)
- [15] A. Levin. Blind motion deblurring using image statistics. In *NIPS*, 2006. [2](#)
- [16] S. Peleg, D. Keren, and L. Schweitzer. Improving image resolution using sub-pixel motion. *Pattern Recognition Letters*, 5:223–226, Mar. 1987. [2](#)
- [17] R. Raskar, A. Agrawal, and J. Tumblin. Coded exposure photography: motion deblurring using fluttered shutter. *ACM Trans. Graph.*, 25(3):795–804, 2006. [2](#), [4](#), [5](#)
- [18] W. Richardson. Bayesian-based iterative method of image restoration. *J. Opt. Soc. of Am.*, 62(1):55–59, January 1972. [2](#)
- [19] R. Schultz and R. Stevenson. Extraction of high-resolution frames from video sequences. *IEEE Trans. Image Processing*, 5:996–1011, June 1996. [2](#)
- [20] G. K. Skinner. X-Ray Imaging with Coded Masks. *Scientific American*, 259:84, Aug. 1988. [2](#)
- [21] J. Sun, N. Zheng, H. Tao, and H. Shum. Image hallucination with primal sketch priors. In *Proc. Conf. Computer Vision and Pattern Recognition*, volume 2, pages 729–736, 2003. [2](#)
- [22] M. Tanaka and M. Okutomi. Theoretical analysis on reconstruction-based super-resolution for an arbitrary psf. In *Proc. Conf. Computer Vision and Pattern Recognition*, volume 2, pages 947–954, 2005. [1](#), [4](#)
- [23] M. Tappen, B. Russell, and W. Freeman. Exploiting the sparse derivative prior for super-resolution and image demosaicing. In *SCTV*, 2003. [2](#)
- [24] H. Ur and D. Gross. Improved resolution from subpixel shifted pictures. In *CVGIP: Graphical Models and Image Processing*, volume 54, pages 181–186, 1992. [2](#)
- [25] A. Zomet and S. Peleg. Efficient super-resolution and applications to mosaics. In *Proc. Int'l Conf. Pattern Recognition*, volume 1, pages 579–583, 2000. [2](#)