

Super-Resolution with Structured Motion

Gabby Litterio, Juan-David Lizarazo-Ferro, Pedro Felzenszwalb, Rashid Zia

Abstract—We consider the limits of super-resolution using imaging constraints. Due to various theoretical and practical limitations, reconstruction-based methods have been largely restricted to small increases in resolution. In addition, motion-blur is usually seen as a nuisance that impedes super-resolution. We show that by using high-precision motion information, sparse image priors, and convex optimization, it is possible to increase resolution by large factors. A key operation in super-resolution is deconvolution with a box. In general, convolution with a box is not invertible. However, we obtain perfect reconstructions of sparse signals using convex optimization. We also show that motion blur can be helpful for super-resolution. We demonstrate that using pseudo-random motion it is possible to reconstruct a high-resolution target using a single low-resolution image. We present numerical experiments with simulated data and results with real data captured by a camera mounted on a computer controlled stage.

Index Terms—Computational Photography, Super-resolution, Compressed Sensing, Deconvolution



1 INTRODUCTION

WE consider the super-resolution problem where the goal is to reconstruct a high-resolution image from one or several low-resolution images. We focus primarily on the use of imaging constraints and low-level image priors, and on the use of motion to capture one or more pictures.

One of our motivations is to implement a compressed sensing system for super-resolution. Using a high-resolution CCD/CMOS sensor, it may be possible to reconstruct an extremely high-resolution image using an appropriate optical and computational imaging system.

Classical reconstruction-based methods relate a collection of low-resolution images to a high-resolution image using a model of the imaging process. Such methods have been largely restricted to small increases in resolution due to theoretical and practical limitations [1], [2]. In addition, motion blur is usually seen as a nuisance that impedes super-resolution [3].

We show that by using multiple static images acquired in a grid of locations and sparse image priors, it is possible to increase resolution by large factors. We note a connection between super-resolution and deconvolution with a box filter, which helps to understand the ambiguities in the super-resolution problem. In particular, the low-resolution images determine most of the Fourier components of the super-resolution image and, in this setting, sparse images can be recovered using convex optimization.

We also show that by using pseudorandom motion it may be possible to reconstruct a high-resolution target using a *single* low-resolution image. Figure 1 illustrates this idea. In this case, we capture an image while the camera/sensor undergoes a large vibration. By solving a convex optimization problem with the imaging constraints and sparsity prior, we obtain an essentially perfect reconstruction of the high-resolution target from a single low-resolution image.

Super-resolution imaging systems have many potential applications. In some settings, there are practical and physical limitations that lead to low-resolution sensors. For the case of visible light photography with high-resolution sensors, our methods could be used to produce gigapixel images using a small camera. High-resolution imaging has many applications in science and in archival/preservation work (including for art, documents, and biological samples). Furthermore, there are many applications where a camera might be constantly moving, such as in aerial imaging.

We implemented a system to demonstrate some of our ideas using a commonly available camera and computer controlled motion stage. The experiments in Section 6 demonstrate results using both static images and images taken with the camera moving at constant velocity.

Throughout the paper, we focus on cases where the effective resolution is limited by the sensor and not by the optical path (diffraction). Even with high-resolution sensors, super-resolution can be useful for obtaining a wider field of view with an appropriate lens. The interlacing operation described in Section 6.2 is reminiscent of a panorama, but interlaces a set of pictures taken with small motion instead of tiling pictures taken with large motion, which leads to a more compact design for a high-resolution camera. The Wiener filter approach for reconstruction is very efficient and yields a significant improvement in resolution in our experiments. In practice, large images can also be processed in parallel via tiling.

We assume throughout the paper that the sensor motion is known, either because it is carefully planned and executed by a high-accuracy actuator or because we have a high-accuracy readout of the sensor position while it is being moved through some other means (such as vibration). For the experiments in Section 6.3, we estimate the camera trajectory during an exposure using a landmark.

2 RELATED WORK

The super-resolution problem has a long history (see, e.g. [1], [2], [4], [5]). From a theoretical point of view, the results

• School of Engineering (Litterio, Felzenszwalb, Zia) and Physics Department (Lizarazo-Ferro, Zia), Brown University, Providence, RI, 02906.
E-mail: gabby_litterio@brown.edu, juan_lizarazoferro@brown.edu, pedro_felzenszwalb@brown.edu, rashid_zia@brown.edu




Fig. 1: Using motion blur for super-resolution.

in [1] and [2] are largely negative, showing that ambiguity in the super-resolution problem grows quickly with the desired increase in resolution. In contrast, here we note that with a particular sampling scheme this ambiguity is concentrated on relatively few Fourier coefficients. Previous work in compressed sensing [6], [7] has shown that sparse image priors can be used to resolve such ambiguities.

[3] proposed the use of a “jitter camera” to improve the resolution of video by quickly moving the sensor in the imaging plane of a camera without introducing motion blur, which was believed to impede super-resolution. Here we argue that motion blur can in fact aid in super-resolution.

Intuitively, the combination of motion blur and discrete sampling can be used to define a compressed sensing system. The approach is related to the use of random filters for compressed sensing [8], [9].

A variety of unique imaging system designs have been developed through the use of compressed sensing methods. Some physical implementations include [10], [11], [12].

Recently, some imaging schemes have been proposed, not to achieve super-resolution, but to remove motion-related artifacts. For example, parabolic lens motion [13], induced by constant acceleration, generates a PSF that is invariant to object velocity. Additionally, coded blur generated by fluttering a camera shutter [14], [15] has been shown to help in removing motion blur. This leads to a surprising idea that it may be beneficial to extend an exposure with a coded shutter to compensate for motion blur.

Many super-resolution methods use training data or other sources of information to improve image resolution. The idea of exemplar based super-resolution goes back to the work in [1] and was implemented at a large scale in [16]. Generic image priors based on self-similarity have also been used for super-resolution [17]. More recent methods have used deep neural networks for super-resolution from a single image [18].

In this paper, we focus on the use of total variation for reconstruction. In practice, alternative image priors, such as ones based on self-similarity or deep networks, could also be implemented using Plug-and-Play methods [19].

Neural models have been increasingly used for image reconstruction, including super-resolution [20]. The use of a neural radiance field (NeRF) allows for the estimation of a continuous image instead of selecting a particular target resolution. In this case, one can implement the imaging model

described here by Monte Carlo integration via sampling a finite number of locations in the imaging plane.

We consider only grayscale images to simplify the presentation and experimental design. For the case of color images captured with a color filter array, one could reconstruct a super-resolution image for each color channel separately, using measurements defined by pixels of that color in the filter array. This is related to the recent approach in [21] which performs super-resolution without explicit demosaicing. By using controlled motion, one can also collect measurements for every pixel using each color in the filter array.

3 IMAGING MODEL

Our goal is to reconstruct a high-resolution image from one or several low-resolution images. In the mathematical model we consider here, the low-resolution images are captured with a sensor that can translate in the imaging plane of a camera. In one setting, we capture a series of images in a grid of high-resolution (sub-pixel) locations. In another setting, we capture one or more images while the sensor is moving; for example, we can record a single or sequence of pictures while vibrating the sensor or moving along a linear trajectory (e.g. as may be the case in satellite imaging).

There are many equivalent practical situations. In the experiments in Section 6, we use a standard camera on a moving stage, which is equivalent to moving the sensor when viewing a planar object parallel to the camera. In a microscope, instead of moving the sensor, we can precisely move the sample using a piezo nanopositioning stage.

Let $g : \mathbb{R}^2 \rightarrow \mathbb{R}$ denote a continuous brightness function in the imaging plane of a camera. We do not model the optics of the camera and focus only on measuring the image projected onto the imaging plane.

An $n \times m$ sensor integrates g over square pixels tiling a rectangular area. We use Δ to denote the side length of a pixel in the sensor. Each pixel integrates the light that falls in a particular square region over a unit of time.

Ignoring measurement noise, if we place the sensor in the imaging plane with bottom-left corner at (x_0, y_0) we capture a discrete image I with

$$I[i, j] = \int_0^\Delta \int_0^\Delta g(x_0 + i\Delta + x, y_0 + j\Delta + y) dx dy.$$

We will also consider the case where the sensor moves while the image is captured. If the sensor follows a trajectory $(x(t), y(t))$, we capture an image I with

$$I[i, j] = \int_0^1 \int_0^\Delta \int_0^\Delta g(x(t) + i\Delta + x, y(t) + j\Delta + y) dx dy dt.$$

Note that translating the sensor by z is equivalent to translating g by $-z$. Let $q(x, y)$ denote the occupancy map defined by the trajectory $(-x(t), -y(t))$. The image I captured by the moving sensor is equal the image of $g \otimes q$ captured by a static sensor.

We consider methods that increase the resolution of the camera sensor to generate a super-resolution image J that would have been recorded by a static sensor with pixels of size Δ/f for an integer factor f defining the resolution increase. This virtual sensor covers an area that is at least as big as the camera sensor and the image J has at least $fn \times fm$ pixels (the area may be larger due to the sensor motion). In our experiments, we have used super-resolution factors as high as $f = 8$.

4 GRID OF IMAGES

Let J be an image generated by a (virtual) sensor with pixel size Δ/f at the origin. Let I be an image captured by a (physical) sensor with pixel size Δ at (x_0, y_0) where

$$x_0 = k(\Delta/f),$$

$$y_0 = l(\Delta/f),$$

with k and l integers. In this case, the low-resolution pixels in I are aligned with the high-resolution pixels in J . In particular, each value in I is the sum of f^2 values in J ,

$$I[i, j] = \sum_{i'=0}^{f-1} \sum_{j'=0}^{f-1} J[k + fi + i', l + fj + j'].$$

We see that I is a decimation (subsampling) of $J \otimes B$ where B is a two-dimensional box filter of width f .

Suppose we capture f^2 different low-resolution images $\{I_{k,l}\}$ by translating the sensor over an $f \times f$ grid of high-resolution locations,

$$L = \{(k(\Delta/f), l(\Delta/f)) \mid 0 \leq k \leq f-1, 0 \leq l \leq f-1\}.$$

Capturing such images involves moving the sensor to sub-pixel locations defined with a step size Δ/f . The total vertical and horizontal displacement required to capture all images is less than Δ .

4.1 Interlacing

The f^2 different images $\{I_{k,l}\}$ can be rearranged into a single $fn \times fm$ image H such that $H = J \otimes B$, where B is a box of width f . Since each image $I_{k,l}$ is a decimation of $J \otimes B$, the rearrangement interlaces the captured images:

$$H[k + if, l + jf] = I_{k,l}[i, j].$$

Figure 2 illustrates the process where we capture multiple low-resolution images and interlace the result to obtain a single “blurry” image H . The interlaced image H has $fn \times fm$ pixels and defines the imaging constraints on J through the relation $H = J \otimes B$.

4.2 Deconvolution with a Box

Recovering the super-resolution image J is equivalent to a deconvolution of H with a two-dimensional box filter of width f . That is, we need to de-blur H to recover J .

Note that a box filter is not a good smoothing or low-pass filter. Convolution with a box yields a signal that is visually blurry but still has quite a lot of high-frequency information. In Figure 3(a) we compare the Fourier transform of a one-dimensional box and a Gaussian of similar width. The Fourier transform of a Gaussian is a Gaussian and the Fourier transform of a box is a sinc. In particular, the magnitude of the Fourier transform of the box decays much slower than the magnitude of the Fourier transform of a Gaussian.

For the case of a two-dimensional discrete (periodic) signal, we illustrate the magnitude of the DFT of B when $f = 4$ in Figure 3(b). Note that there are relatively few coefficients in the DFT of B with magnitude close to 0. We can therefore directly recover most of the coefficients in the Fourier transform of J from the Fourier transform of H , except for isolated frequencies where the Fourier transform of B is zero or close to zero.

Previous work [1], [2] showed that the ambiguity in J grows exponentially with f . We note, however, that even for large factors f , the constraints defined by $J = H \otimes B$ identify *most* of the Fourier coefficients of J .

The number of times the Fourier transform of a discrete periodic box of width f crosses zero is $O(f)$. Therefore, the ambiguity in J is concentrated on relatively few Fourier coefficients. In this setting, sparse image priors and convex optimization are known to yield surprisingly good reconstructions [6], [7].

For the specific case of deconvolution with a box, we have shown [22] that convex optimization using ℓ_1 regularization can lead to perfect reconstruction of sparse one-dimensional signals.

Figure 4 illustrates a numerical experiment with a simulated camera. The deconvolution of H using a total variation prior yields an almost perfect reconstruction including very high-resolution details.

4.3 Wiener Deconvolution

Using a Gaussian prior on the Fourier coefficients of J leads to a fast Wiener deconvolution algorithm via the FFT and an inverse transform. The result corresponds to a MAP estimate of the super-resolution image J . The prior is defined by a single parameter γ used to shrink the estimated Fourier coefficients.

$$\bar{B} = \text{FFT}(B)$$

$$\bar{H} = \text{FFT}(H)$$

$$\bar{J} = \frac{\bar{H}}{\bar{B} + \gamma}$$

$$J = \text{IFFT}(\bar{J})$$

4.4 Sparse Reconstruction

As mentioned above, previous work has shown that sparse signals can often be recovered from incomplete frequency information [6], [7].




Fig. 2: Interlacing $f \times f$ low-resolution images taken in a grid of sub-pixel locations, we obtain the convolution of the high-resolution target with a two-dimensional box filter of width f .




Fig. 3: (a) One-dimensional box and Gaussian (top) and their Fourier transforms (bottom). The red plots correspond to a Gaussian and the blue plots correspond to a box. (b) Bode plot (magnitude of Fourier transform in dB) of a two-dimensional discrete box filter of width 4.

In practice, we regularize the reconstruction of the super-resolution image J using the total variation (TV) measure, which leads to a convex optimization problem,

$$\min_J \|H - (J \otimes B)\|_2^2 + \lambda TV(J)$$

where $TV(J)$ is the total-variation of J ,

$$TV(J) = \sum_{i,j} |J[i, j] - J[i + 1, j]| + |J[i, j] - J[i, j + 1]|.$$

and λ is a regularization parameter.

5 SUPER-RESOLUTION BY DEBLURRING

In the previous section, we considered the case of multiple static images taken from different sensor locations. Here we consider the case where one or more images are captured while the sensor is moving.

First, we note that a small amount of motion blur can be helpful for super-resolution. Consider the problem of localizing a point source in one dimension; without motion, we can only localize the point source to the pixel size. Now suppose we move the sensor/camera with constant velocity by one pixel while taking a picture. Figure 5 illustrates this situation. The point source is spread over two pixels and we observe two non-zero values $I[k]$ and $I[k + 1]$. Together,




Fig. 4: Super-resolution with $f = 8$. By interlacing 64 low-resolution images, we obtain the measurements $H = J \otimes B$. Deconvolving H with a TV prior leads to an almost perfect reconstruction.




Fig. 5: We can localize a point source to sub-pixel location if the camera undergoes a motion of one pixel during capture. A motion of one pixel spreads the source to two pixels, and the relative values determine the source location. On the left, we have the continuous brightness function and the image captured with a static sensor. On the right, the sensor moves with constant velocity by one pixel during capture.

these measurements determine both the total light intensity and the precise point source location,

$$a = b + c,$$

$$t_0 = (k + 0.5)\Delta + c/(b + c).$$

Optical blurring, such as due to diffraction, can lead to a similar effect (see, e.g. [23]) but the use of motion may have advantages since, as discussed in the previous section, convolution with a box preserves more information when compared to a Gaussian filter of similar width.

We also note that for “sparse” scenes, we can use *large* motions to (in essence) take multiple low-resolution images in different areas of the sensor. Shown below in Figure 6 is a simulation where we vibrate a sensor while capturing an image, which leads to the somewhat counterintuitive notion that shaking the camera during an exposure can lead to sharper pictures.

Being able to take pictures while the sensor/camera is moving allows for faster acquisitions, which is often important. In this case, there is no need to wait for the sensor to move and stabilize in one position before taking a picture; instead, we can follow a continuous trajectory and take a sequence of pictures along the way.

As in the previous section, let J be a super-resolution image generated by a virtual sensor at the origin. Now we let I be an image captured by a moving sensor. Let $(k(t), l(t))$ for $0 \leq t \leq 1$ denote the trajectory of the sensor over the grid of high-resolution pixel locations. That is, at time t , the sensor is at location $(k(t)(\Delta/f), l(t)(\Delta/f))$.

Let Q be the occupancy map of $(-k(t), -l(t))$. That is, $Q[k, l]$ is the total fraction of time this trajectory spends at location (k, l) . When the sensor undergoes a continuous trajectory we use linear interpolation to define a discrete occupancy map Q . Let B be a box of width f . The relationship between J and I is captured by a decimation by factor f of the sequential convolution of J with Q and B . The convolution of J with Q accounts for the sensor motion, while the convolution with B and decimation accounts for

the relative size of the physical and virtual pixels.

Let $(J) \downarrow_f$ denote the decimation of J by a factor f . The imaging constraints are summarized by the linear system,

$$I = (J \otimes Q \otimes B) \downarrow_f.$$

This system of equations has many fewer constraints than variables, but it can lead to perfect reconstructions with an appropriate image prior as we demonstrate in simulations below.

In some applications we may have a sequence of low-resolution images $I_1, \dots, I_s$ taken over one or multiple trajectories. Let O_j denote the occupancy map of the trajectory followed during the acquisition of I_j .

We can reconstruct a single super-resolution image using the data from multiple trajectories and a sparsity prior. To estimate a sparse image use ℓ_1 regularization,

$$\min_J \sum_{k=1}^s \|I_k - (J \otimes Q_k \otimes B) \downarrow_f\|_2^2 + \lambda \|J\|_1.$$

Similarly we can regularize the reconstruction using TV,

$$\min_J \sum_{k=1}^s \|I_k - (J \otimes Q_k \otimes B) \downarrow_f\|_2^2 + \lambda TV(J).$$

Here again λ is a regularization parameter.

5.1 Simulations

We evaluate the quality of the reconstructions obtained using a single exposure when the sensor moves in different trajectories.

Figure 6(a) shows two examples of the high-resolution targets we used for the experiments. Each target is a 128×128 image with several random characters placed in random locations. Figure 6(b) shows the results of taking a low-resolution picture of each target using a 32×32 sensor. In this case, there is no sensor motion. We can see that with such a low-resolution picture, the characters become completely unreadable.

Figure 6(c) shows the results of imaging the two targets from Figure 6(a) with a moving sensor. In each case, the sensor undergoes a large motion. Figure 6(d) shows the reconstructions of the targets from a single low-resolution image, captured while the sensor is moving. In these examples, we obtain a perfect super-resolution image from a single low-resolution image by leveraging motion blur.

For the experiments in this section, we treat the high-resolution target as periodic. The sensor motion therefore leads to a cyclic convolution between the target and the sensor occupancy map. We used the ℓ_1 norm of the target image J to regularize the reconstructions.

We consider the following trajectories for the sensor:

- 1) 1000 random displacements (shifts). In this case, the trajectory is a sequence of 1000 random locations. Each location is selected independently from a uniform distribution in the high-resolution grid.
- 2) Random walk with 1000 steps (walk). In this case, the sensor undergoes a random walk in the grid of high-resolution locations.
- 3) Vibration (vibration). The x and y coordinates of the sensor are defined by two sinusoids.

Figure 7 shows the mean RMS error obtained with each trajectory for targets of increasing complexity. We see that for sparse targets, where the number of characters in the image is small, it is possible to obtain perfect 128×128 images using a single 32×32 blurry picture.

6 EXPERIMENTAL RESULTS

We implemented the deconvolution algorithms in python using numpy for the Wiener deconvolution method and both cvxpy and pylops for computing reconstructions using sparsity priors. We found that using the OSQP solver within cvxpy yields good results with small images, as seen in the simulations throughout the paper. For the experiments with real data, we used the split-bregman algorithm implemented in pylops due to the computational demands of processing larger images.

All of the experiments were done with a desktop computer with a 13th Gen Intel(R) Core(TM) i5-13400F CPU with 16 cores and 32GB of RAM.

When $f = 8$, reconstructing a 1200×1520 image from 64 separate 150×190 low-resolution images with the Wiener deconvolution algorithm took about 0.4 seconds. Computing a TV reconstruction using pylops with the same data took between 5 and 10 minutes depending on various tolerances and algorithm parameters.

6.1 Setup

The data used in the following reconstructions consists of sets of images captured by a monochrome 1.6 megapixel (1440×1080) CMOS camera (Thorlabs Zelux) with $3.45 \mu\text{m}$ square pixels and a wide-angle 8 mm f1.8 Edmund Optics lens opened to maximum aperture and focused at approximately 3 meters. This particular choice of lens and sensor leads to an optical system that is not diffraction limited despite the relatively small pixel size.

The camera was mounted on a movable stage and controlled by a nanopositioning system (Nanomotion II) that uses dc stepper motors to move the stage at nanometer level resolution in two separately controlled axes. Movements made by the stage have high precision and repeatability, though no position readout is available. The stage is connected to and controlled by a computer via a serial port.

A static platform was mounted 3 meters away and parallel to the camera to hold targets to be imaged. The illumination source was a battery powered LED tube light (10 inch PavoTubeII 6C RGBWW).

We imaged a variety of targets including small three-dimensional objects (ex. Ketchup packet and LEGO(R) sets) and multiple $8.5'' \times 11''$ pages of paper containing designs printed by a standard black and white laser printer, such as a page of text and a to-scale version of the 1951 USAF resolution target.

The USAF target is a standard test chart used to determine the resolution of an optical system. It consists of vertical and horizontal sets of three bars at varying sizes and spatial frequencies that are organized in numbered groups, each with six labeled sets. Once imaged, intensity profiles across the bars can be analyzed; past the resolution limit, three distinct peaks/valleys will no longer be visible in




Fig. 6: (a) High-resolution targets with 10 random characters each. (b) Images recorded by a static low-resolution sensor. (c) Sensor trajectory (vibration on the top and random translations/shifts on the bottom). (d) Low-resolution images recorded with a sensor undergoing large motions. (e) Reconstruction of the high-resolution target from the blurry images in (d).




Fig. 7: Mean reconstruction error for targets with different numbers of characters. Each target is 128×128 and we capture a single 32×32 image undergoing motion.

the profile. The group and element number of the smallest resolvable set is used to look up the width and spacing of these bars, and this value, in micrometers or lp/mm, is a measure of the resolution of the system.

For each set of experiments, flat-field images were acquired by imaging a blank sheet of paper, which served as the background for all targets.

6.2 Grid Exposures

In our setup we move the camera instead of the sensor. To acquire a static grid of images, the stage was moved through a series of positions in a plane parallel to the target, coming to a stop at each position and waiting for two frames to be

captured, before continuing. The overall grid motion was repeated five times. Averaging many frames reduces the effects of sensor noise and running multiple trials of the grid sequence minimizes the effects of temporal illumination variation during the acquisition.

The number of acquired low-resolution images is determined by the desired resolution increase factor f . The actual position of these acquired images depends on the magnification factor induced by the camera optics and the distance to the target.

Moving the camera is equivalent to moving the target, and the magnification m of the system relates the distance traveled by the camera to the distance traveled by the image on the sensor. We can determine m using the optical properties of the camera and the distance to the target. In practice, we find m by imaging an object with known length and calculating its size on the sensor by counting the number of pixels it spans and multiplying by the pixel size.

With grid capture (see Section 6.2), the image of the target moves by less than Δ in both the vertical and horizontal directions, with pictures taken at f sub-pixel positions, for a total of $f \times f$ image locations. Taking into account the magnification m , the distance the stage travels between neighboring grid positions, i.e. step size, is $\frac{\Delta}{fm}$.

For $f = 8$, we use 64 grid locations. The magnification of our system was calculated to be 0.002484 (we use a wide-angle lens), which gives a step size of 0.17361 mm.

Due to the large field of view of the camera, our targets occupy a relatively small (150×190) area in the images acquired, which means that for the experiments presented here, we are only using a small region of the physical sensor in the camera. In other applications, the same setup would allow for capturing very wide field of views.

Figure 9 shows several results, each a small ROI of the target area. We note that in each example we can see high-resolution details appear in the super-resolution images (e.g.




Fig. 8: (a) The camera is attached to a linear stage that adjusts the distance between the camera and the target. The linear stage is secured to the Nanomotion II stage. (b) Example of printed design and small object placed on the target platform.

punctuation and to some extent serifs, elements in USAF group -1, the serrated top edge of the ketchup packet).

Figure 10 shows that our magnification model holds for three-dimensional objects that protrude out of the target plane. The front of the LEGO(R) bus is approximately 15 cm away from the calibrated target plane and high-resolution details on the minifigures are still resolved.

6.3 Scan Motion

To acquire images with motion, the stage was moved at constant horizontal velocity over a “long” travel range corresponding to a length of many pixels. As the stage moved, five images were captured with a short random delay set before the start of each acquisition. Once the motion across the row was completed, the height of the stage was increased by sub-pixel steps (same as in grid exposures) and the motion was repeated. The total number of heights corresponds to the selected factor f .

Since our stage is slow, we used a relatively long exposure to capture images with motion blur. The exposure was determined by how long it takes for the image on the sensor to move by Δ (the size of the sensor pixel). The maximum velocity of the stage is only 2.5 mm/s and therefore, for these experiments we used an exposure time of about five hundred milliseconds.

Due to the fact that the stage has no position readout capabilities and the camera cannot be synchronized with it, the initial camera position for each acquisition must be estimated. The printed targets include a black rectangle, which provides a step edge that is used for the position estimation; similar to the blurred point source shown in Figure 5, a sharp edge can be used for localization. The exposure time, vertical position, and velocity of the stage during each capture are known, so once the initial position has been estimated, the trajectory of motion during the acquisition is fully defined.

Figure 11 illustrates some results of the scanning motion procedure. These examples demonstrate the use of a moving camera to increase resolution beyond what is visible in a static image, despite the induced motion blur.

7 CONCLUSION AND FUTURE WORK

We described several methods for increasing the limits of super-resolution with a moving sensor/camera.

Using a grid of static images, the ambiguity in super-resolution is restricted to relatively few Fourier coefficients. This enables fast reconstruction using a Wiener filter, and higher effective super-resolution factors with sparse priors and compressed sensing techniques.

Note that by using two different magnifications during acquisition, or a sensor with multiple pixel sizes, one could potentially remove all of the ambiguity in the super-resolution problem. Figure 12 illustrates the Fourier transform of a discrete periodic box of width 4 and similarly of a box of width 7. We see that the zeros of the two Fourier transforms occur in different places; if we were to record the convolution of a high-resolution image with boxes of two sizes, we would eliminate essentially all ambiguity in the super-resolution problem. Taking pictures while moving at different speeds could also be used to integrate the brightness function over regions of different sizes.

We have also shown that motion blur can in fact be used to improve the resolution with a single image. This leads to the idea that to take a sharper picture one should shake or vibrate the camera. In synthetic experiments, we have shown that simply vibrating the image sensor during capture can lead to higher resolution images. The key here is that we solve a motion deblurring problem at the same time that we increase the image resolution, which can be accomplished as long as we have high-resolution trajectory information.

REFERENCES

- [1] S. Baker and T. Kanade, “Limits on super-resolution and how to break them,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 24, no. 9, pp. 1167–1183, 2002.
- [2] Z. Lin and H.-Y. Shum, “Fundamental limits of reconstruction-based superresolution algorithms under local translation,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 26, no. 1, pp. 83–97, 2004.
- [3] M. Ben-Ezra, A. Zomet, and S. K. Nayar, “Jitter camera: High resolution video from a low resolution detector,” in *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2.




Fig. 9: Super-resolution using an 8x8 grid of static images. We show a small crop of various targets.

- [4] M. Irani and S. Peleg, "Super resolution from image sequences," in *Proceedings. 10th International Conference on Pattern Recognition*, vol. 2. IEEE, 1990, pp. 115–120.
- [5] D. Capel and A. Zisserman, "Super-resolution enhancement of text image sequences," in *Proceedings: 15th International Conference on Pattern Recognition*. IEEE, 2000.
- [6] E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Transactions on information theory*, vol. 52, no. 2, pp. 489–509, 2006.
- [7] D. L. Donoho and P. B. Stark, "Uncertainty principles and signal recovery," *SIAM Journal on Applied Mathematics*, vol. 49, no. 3, pp. 906–931, 1989.
- [8] J. Romberg, "Compressive sensing by random convolution," *SIAM Journal on Imaging Sciences*, vol. 2, no. 4, pp. 1098–1128, 2009.
- [9] J. A. Tropp, M. B. Wakin, M. F. Duarte, D. Baron, and R. G. Baraniuk, "Random filters for compressive sampling and reconstruction," in *2006 IEEE International Conference on Acoustics Speech and Signal Processing Proceedings*, vol. 3.
- [10] R. Fergus, A. Torralba, and W. T. Freeman, "Random lens imaging," Massachusetts Institute of Technology Computer Science and Artificial Intelligence Laboratory, Tech. Rep., 2006.
- [11] N. Antipa, G. Kuo, R. Heckel, B. Mildenhall, E. Bostan, R. Ng, and L. Waller, "DiffuserCam: lensless single-exposure 3d imaging," *Optica*, vol. 5, no. 1, pp. 1–9, 2017.
- [12] M. F. Duarte, M. A. Davenport, D. Takhar, J. N. Laska, T. Sun, K. F. Kelly, and R. G. Baraniuk, "Single-pixel imaging via compressive sampling," *IEEE signal processing magazine*, vol. 25, no. 2, pp. 83–91, 2008.
- [13] A. Levin, P. Sand, T. S. Cho, F. Durand, and W. T. Freeman, "Motion-invariant photography," *ACM Transactions on Graphics*, vol. 27, no. 3, pp. 1–9, 2008.
- [14] R. Raskar, A. Agrawal, and J. Tumblin, "Coded exposure photography: motion deblurring using fluttered shutter," *ACM Transactions on Graphics*, vol. 25, no. 3, pp. 795–804, 2006.
- [15] Y. Tendo, J.-M. Morel, and B. Rougé, "The flutter shutter paradox," *SIAM Journal on Imaging Sciences*, vol. 6, no. 2, pp. 813–847, 2013.
- [16] L. Sun and J. Hays, "Super-resolution from internet-scale scene matching," in *2012 IEEE International conference on computational photography*, pp. 1–12.
- [17] D. Glasner, S. Bagon, and M. Irani, "Super-resolution from a single image," in *12th international conference on computer vision*. IEEE, 2009, pp. 349–356.



(a) Static image

(b) Wiener reconstruction

(c) TV reconstruction

Fig. 10: Super-resolution using an 8x8 grid of static images. The bottom row shows a crop of the front section of the bus taken from the full image shown in the top row. Using a moving camera instead of a moving sensor requires the camera motion to be proportional to the distance from the object. Nonetheless we can image three-dimensional objects that are relatively close to a target plane. In this case, a constant magnification model is still accurate.

- [18] C. Dong, C. C. Loy, K. He, and X. Tang, "Image super-resolution using deep convolutional networks," *IEEE transactions on pattern analysis and machine intelligence*, vol. 38, no. 2, pp. 295–307, 2015.
- [19] U. S. Kamilov, C. A. Bouman, G. T. Buzzard, and B. Wohlberg, "Plug-and-play methods for integrating physical and learned models in computational imaging: Theory, algorithms, and applications," *IEEE Signal Processing Magazine*, vol. 40, no. 1, pp. 85–97, 2023.
- [20] Y. Xie, T. Takikawa, S. Saito, O. Litany, S. Yan, N. Khan, F. Tombari, J. Tompkin, V. Sitzmann, and S. Sridhar, "Neural fields in visual computing and beyond," in *Computer Graphics Forum*, vol. 41. Wiley Online Library, 2022, pp. 641–676, issue: 2.
- [21] B. Wronski, I. Garcia-Dorado, M. Ernst, D. Kelly, M. Krainin, C.-K. Liang, M. Levoy, and P. Milanfar, "Handheld multi-frame super-resolution," *ACM Transactions on Graphics*, vol. 38, no. 4, pp. 1–18, 2019.
- [22] P. Felzenszwalb, "Deconvolution with a box," Tech. Rep. arXiv:2407.11685, 2024.
- [23] G. Schiebinger, E. Robeva, and B. Recht, "Superresolution without separation," *Information and Inference: A Journal of the IMA*, vol. 7, no. 1, pp. 1–30, 2018.




Fig. 11: Super-resolution using 4 horizontal scans with constant velocity and 4 exposures per scan. (a) shows an ROI from a static image. (b) shows the super-resolution result using 16 images taken with constant velocity motion.




Fig. 12: Bode plots of a two-dimensional box filter of width 4 (a), width 7 (b), and the sum of the two (c). The last plot illustrates how using pixels of two different sizes should eliminate most of the ambiguity for super-resolution.