

Jitter Camera: High Resolution Video from a Low Resolution Detector *

Moshe Ben-Ezra, Assaf Zomet, and Shree K. Nayar
Computer Science Department
Columbia University
New York, NY, 10027
E-mail: {moshe, zomet, nayar}@cs.columbia.edu

Abstract

Video cameras must produce images at a reasonable frame-rate and with a reasonable depth of field. These requirements impose fundamental physical limits on the spatial resolution of the image detector. As a result, current cameras produce videos with a very low resolution. The resolution of videos can be computationally enhanced by moving the camera and applying super-resolution reconstruction algorithms. However, a moving camera introduces motion blur, which limits super-resolution quality. We analyze this effect and derive a theoretical result showing that motion blur has a substantial degrading effect on the performance of super resolution. The conclusion is, that in order to achieve the highest resolution, motion blur should be avoided.

Motion blur can be minimized by sampling the space-time volume of the video in a specific manner. We have developed a novel camera, called the "jitter camera," that achieves this sampling. By applying an adaptive super-resolution algorithm to the video produced by the jitter camera, we show that resolution can be notably enhanced for stationary or slowly moving objects, while it is improved slightly or left unchanged for objects with fast and complex motions. The end result is a video that has a significantly higher resolution than the captured one.

1. Why is High-Resolution Video Hard?

Improving the spatial resolution of a video camera is different from doing so with a still camera. Merely increasing the number of pixels of the detector reduces the amount of light received by each pixel, and hence increases the noise. With still images, this can be overcome by prolonging the exposure time. In the case of video, however, the exposure time is limited by the desired frame-rate. The amount of light incident on the detector can also be increased by widening the aperture, but with a significant re-

*This research was conducted at the Columbia Vision and Graphics Center in the Computer Science Department at Columbia University. It was funded in parts by an ONR Contract (N00014-03-1-0023) and an NSF ITR Grant(IIS-00-85864).

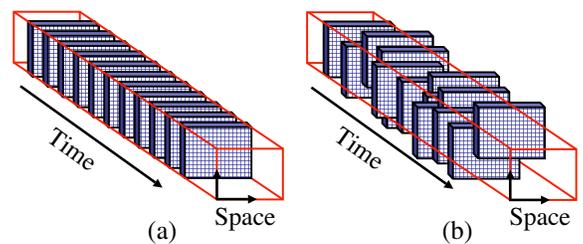


Figure 1: Conventional video cameras sample the continuous space-time volume at regular time intervals and fixed spatial grid locations as shown in (a). The space-time volume can be sampled differently, for example, by varying the location of the sampling grid as shown in (b) to increase the resolution of the video. A moving video only approximates (b) due to motion blur.

duction of the depth of field. The spatial resolution of a video detector is therefore limited by the noise level of the detector, the frame-rate (temporal resolution) and the required depth of field¹. Our purpose is to make a judicious use of a given detector, that will allow a substantial increase of the video resolution by a resolution-enhancement algorithm.

Figure 1 shows a *continuous* space-time video volume. A slice of this volume at a given time instance corresponds to the image appearing on the image plane of the camera at this time. This volume is sampled both spatially and temporally, where each pixel integrates light over time and space. Conventional video cameras sample the volume in a simple way, as shown in Figure 1(a), with a regular 2D grid of pixels integrating over regular temporal intervals and at fixed spatial locations. An alternative sampling of the space-time volume is shown in Figure 1(b). The 2D grid of pixels integrates over the same temporal intervals, but at different spatial locations. Given a 2D image detector, how should we sample the space-time volume to obtain the highest spatial resolution²?

¹The optical transfer function of the lens also imposes a limit on resolution. In this paper we ignore this limit as it is several orders of magnitudes above the current resolution of video.

²Increasing the temporal resolution [18] is not addressed in this paper.

There is a large body of work on resolution enhancement by varying spatial sampling, commonly known as super-resolution reconstruction [5, 6, 8, 10, 14, 17]. Super-resolution algorithms typically assume that a set of displaced images are given as input. With a video camera, this can be achieved by moving the camera while capturing the video. However, the camera's motion introduces motion blur. This is a key point in this paper: in order to use super-resolution with a conventional video camera, the camera must move, but when the camera moves, it introduces motion blur which reduces resolution.

It is well known that an accurate estimation of the motion blur parameters is non-trivial, and requires strong assumptions about the camera motion during integration [2, 14, 15, 19]. In this paper, we show that *even when an accurate estimate of the motion blur parameters is available*, motion blur has a significant influence on the super-resolution result. We derive a *theoretical lower bound*, indicating that the expected performance of any super-resolution reconstruction algorithm deteriorates as a function of the motion blur magnitude. The conclusion is that, in order to achieve the highest resolution, motion blur should be avoided.

To achieve this, we propose the "jitter camera," a novel video camera that samples the space-time volume at different locations without introducing motion blur. This is done by instantaneously shifting the detector (e.g. CCD) between temporal integration periods, rather than continuously moving the entire video camera during the integration periods. We have built a jitter camera, and developed an adaptive super-resolution algorithm to handle complex scenes containing multiple moving objects. By applying the algorithm to the video produced by the jitter camera, we show that resolution can be enhanced significantly for stationary or slowly moving objects, while it is improved slightly or left unchanged for objects with fast and complex motions. The end result is a video that has higher resolution than the captured one.

2. How Bad is Motion Blur for Super-resolution?

The influence of motion blur on super resolution is well understood when all input images undergo the same motion blur [1, 11]. It becomes more complex when the input images undergo *different* motion blurs, and details that appear blurred in one image, appear sharp in another image. We address the influence of motion blur for any combination of blur orientations.

Super-resolution algorithms estimate the high resolution image by modeling and inverting the imaging process. Analyzing the influence of motion blur requires a definition for super-resolution "hardness" or the "invertibility" of the imaging process. We use a linear model for the imaging process [1, 8, 10, 14], where the intensity of a pixel in the input image is presented as a linear combination of the intensities in the unknown high resolution image:

$$\vec{y} = A\vec{x} + \vec{z}, \quad (1)$$

where \vec{x} is a vectorization of the unknown discrete high resolution image, \vec{y} is a vectorization of all the input images, and the

imaging matrix A encapsulates the camera displacements, blur and decimation [8]. The random variable \vec{z} represents the uncertainty in the measurements due to noise, quantization error and model inaccuracies.

Baker and Kanade [1] addressed the invertibility of the imaging process in a noise-free scenario, where \vec{z} represents the quantization error. In this case, each quantized input pixel defines two inequality constraints on the super-resolution solution. The combination of constraints forms a *volume of solutions* that satisfy all quantization constraints. Baker and Kanade suggest to use the *volume of solutions* as a measure of uncertainty in the super resolution solution. Their paper shows the benefits in measuring the *volume of solutions* over the standard matrix conditioning analysis.

We measure the influence of motion blur by the volume of solutions. To keep the analysis simple, the following assumptions are made. First, the optical blur and motion blur are shift-invariant. Second, the input images are related geometrically by a 2D translation. Third, the number of input pixels equals the number of output pixels. Under the last assumption, the dimensionality n^2 of \vec{x} equals the dimensionality of \vec{y} . Since the uncertainty due to quantization is an n^2 -dimensional unit cube, the volume of solutions for a given imaging matrix A is computed from the absolute value of its determinant

$$\text{vol}(A) = \left| \frac{1}{|A|} \right|. \quad (2)$$

In [3], we derive a simplified expression for $\text{vol}(A)$ as a function of the imaging parameters. We define a function $s(A)$ that encapsulates the influence of motion blur on the volume of solutions. This function depends only on the motion blur and the camera displacements, and satisfies:

$$\text{vol}(A) \propto s(A)^{n^2}$$

$s(A)$ is therefore a relative measure for the volume of solutions that is independent of the optical blur and detector's integration function, and is normalized to account for the image size n^2 . A detailed expression for $s(A)$ appears in [3].

We used $s(A)$ to perform a quantitative analysis of the effect of motion blur on the volume of solutions, as shown in figure 2. The value of $s(A)$ was computed for 4 input images undergoing motion blurs caused by constant-velocity motions. Different input images may have different motion blur orientations. Specifically, let \vec{l}_j be a vector describing the motion blur trajectory for the j -th input image: During integration, the projected image moves at a constant velocity from $-\frac{\vec{l}_j}{2}$ to $\frac{\vec{l}_j}{2}$. Each graph in figure 2 shows the value of $s(A)$ as a function of the length of the four motion blur trajectories $\{\|\vec{l}_j\|\}_{j=0}^3$. The different graphs correspond to different configurations of blur orientations in four input images. The camera displacements were taken to be the optimal ones [3], corresponding to the high resolution grid.

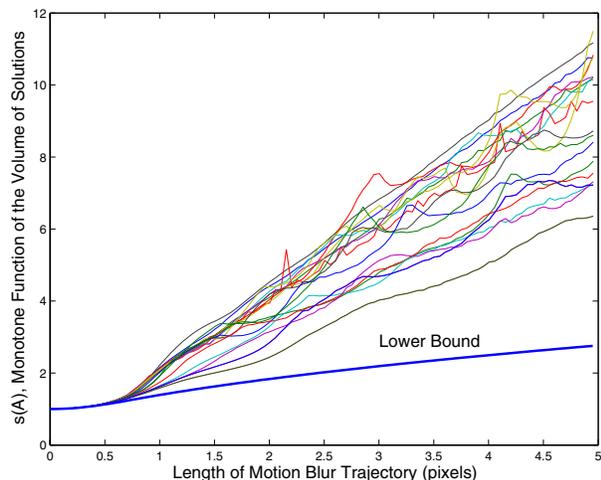


Figure 2: We measure the super-resolution “hardness” by the volume of plausible high-resolution solutions [1]. The volume of solutions is proportional to $s(A)^{n^2}$ where n^2 is the high resolution image size. The graphs show the value of $s(A)$ as a function of the length of the motion-blur trajectories $\{\|l_j\|\}_{j=0}^3$. We show a large number of graphs computed for different configurations of blur orientations. The thick graph (blue line) is the lower bound of $s(A)$, for any combination of motion blur orientations. In all shown configurations, the motion blur has a significant influence on $s(A)$ and hence on the volume of solutions. The increase in the volume of solutions can explain the increase in reconstruction error in super-resolution shown in Figure 3.

It can be seen that in **all** selected motion blur configurations $s(A) \propto vol(A)^{\frac{1}{n^2}}$ increases as a function of the length of the motion blur trajectories $\{\|l_j\|\}$. The thick blue line is the lower bound of $s(A)$, whose derivation can be found in [3]. This bound holds for any configuration of blur orientations.

The findings above confirm that, at least for our assumptions, *any motion blur is bad for super-resolution and the larger the motion blur, the larger the volume of solutions.*

Figure 3(a) shows super-resolution results of simulated inputs with and without motion blur. Motion blur as small as 3.5 pixels degrades the super-resolution result, such that some of the letters are unreadable. Figure 3(b) presents the RMS error in the reconstructed super-resolution image as a function of the extent of the known motion blur. It can be seen that the RMS error increases as a function of the motion blur magnitude. This effect is consistent with the theoretical observations made above.

3. Jitter Video: Sampling without Motion Blur

Our analysis showed that sampling with minimal motion blur is important for super-resolution. Little can be done to prevent motion blur when the camera is moving³ or when objects in the scene are moving. Therefore, our main goal is to sample at different spatial locations while avoiding motion blur in static

³Small camera shakes can be eliminated by *optical* lens stabilization systems, which stabilize the image before it is integrated.

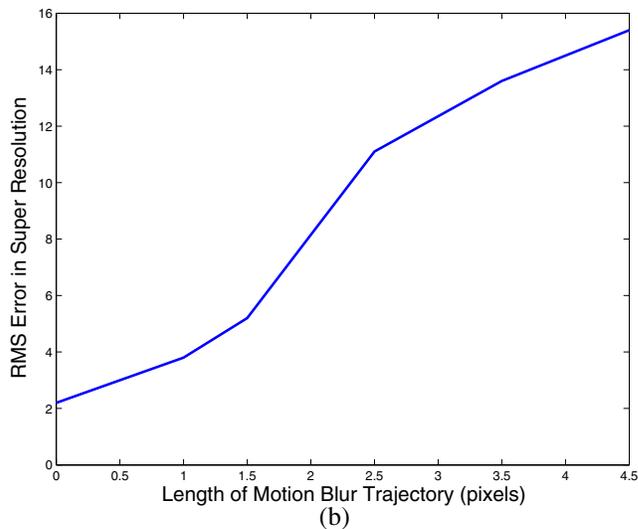
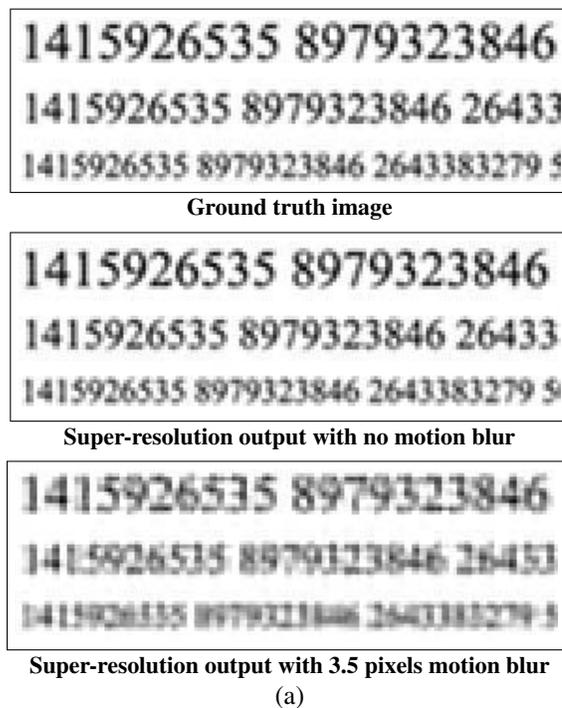


Figure 3: The effect of motion blur on super-resolution with a known *simulated* motion blur. (a) The top image is the original ground-truth image. The middle image is the super-resolution result for 4 simulated input images with no motion blur. This image is almost identical to the ground truth image. The bottom image is a super-resolution result for 4 *simulated* input images with motion blur of 3.5 pixels. Two images with horizontal blur and two with vertical blur were used. The algorithm used the *known* simulated motion blur kernels and the known displacements. The degradation in the super-resolution result due to motion blur is clearly visible. (b) The graph shows the grey level RMS error in the super-resolution image as a function of motion blur trajectory length.

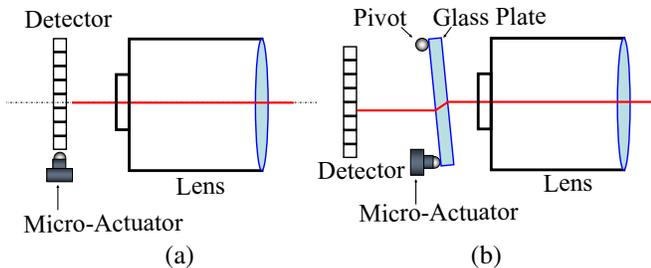


Figure 4: A jitter video camera shifts the sampling grid accurately and instantaneously. This can be achieved using micro-actuators, which are both fast and accurate. The actuator can shift the detector as shown in (a), or it can be used to operate a simple optical device, such as the tilted glass plate shown in (b), in order to optically move the image with respect to the static detector.

regions of the image.

The key to avoiding motion blur is *synchronous* and *instantaneous* shifts of the sampling grid between temporal integration periods, rather than a continuous motion during the integration periods. In [3] we show that the volume of solutions can be minimized by properly selecting the grid displacements. For example, in the case of four input images, one set of optimal displacements is achieved by shifting the sampling grid by half a pixel horizontally and vertically. Implementing these abrupt shifts by moving a standard video camera with a variable magnification factor is non-trivial⁴. Hence we propose to implement the shifts of the sampling grid inside the camera.

Figure 4 shows two possible ways to shift the sampling grid instantaneously. Figure 4(a) shows a purely mechanical design, where the detector (e.g. CCD) is shifted by actuators to change the sampling grid location. If the actuators are fast and are activated *synchronously* with the reading cycle of the detector, then the acquired image will have no motion blur due to the shift of the detector. Figure 4(b) shows a mechanical-optical design. A flat thin glass plate is used to shift the image over the detector. An angular change of a 1mm thick plate by one degree shifts the image by $5.8\mu m$, which is of the order of a pixel size. Since the displacement is very small relative to the focal length, the change of the optical path length results with negligible effect on the focus (the point spread area is much smaller than the area of a pixel). The mechanical-optical design shown Figure 4(b) has been used for high-resolution still-imaging, for example by Pixera [7], where video related issues such as motion blur and dynamic scenes do not arise.

An important point to consider in the design of a jitter camera is the quality of the camera lens. With standard video cameras, the lens-detector pair is matched to reduce spatial aliasing in the detector. For a given detector, the matching lens attenuates

⁴A small uniform image displacement can be approximated by rotating the camera about the X, Y axes. However, the rotation extent depends on the exact magnification factor of the camera, which is hard to obtain. In addition, due to camera's mass, abrupt shifting of the camera is challenging.

the spatial frequencies higher than the Nyquist frequency of the detector. For a jitter camera, higher frequencies are useful since they are exploited in the extraction of the high resolution video. Hence, the selected lens should match a detector with a higher (the desired) spatial resolution.

4. The Jitter Camera Prototype

To test our approach, we have built the jitter camera prototype shown in Figure 5. This camera was built using a standard 16mm television lens, a Point-Grey [16] Dragon-Fly board camera, and two Physik Instrumente [9] micro-actuators. The micro-actuators and the board camera were controlled and synchronized by a Physik Instrumente Mercury stand-alone controllers (not shown).

The jitter camera is connected to a computer using a standard firewire interface, and therefore it appears to be a regular firewire camera.

We used in our prototype two DC-motor actuators, which enable a frame-rate of approximately 8 frames per second. Newly developed piezoelectric based actuators can offer much higher speed than DC-motor based actuators. Such actuators are already used for *camera shake compensation* by Minolta [13], however they are less convenient for prototyping at this point in time.

The camera operates as follows:

1. At power up the actuators are moved to a fixed home-position.
2. For each sampling position in $[(0,0),(0,0.5),(0.5,0.5),(0.5,0)]$ pixels do
 - Move the actuators to the next sampling position.
 - Bring the actuators to a full stop.
 - Send a trigger signal to the camera to initiate frame integration and wait during integration duration.
 - When the frame is ready, the camera sends it to the computer over the Firewire interface.
3. End loop
4. Repeat process from step (2).

To evaluate the accuracy of the jitter mechanism we captured a sequence of images with the jitter camera, computed the motion between frames to a sub-pixel accuracy [4] and compared the computed motion to the expected value. The results are shown in Figure 6. The green circles show the expected displacements, and the red diamonds show the actual displacements over multiple cycles. We can see that the accuracy of the jitter mechanism was better than 0.1 pixel. We can also see that while some error is accumulated along the path, the camera accurately returns to its zero position, thus preventing drift.

The resolution of the computed high-resolution video was 1280×960 , which has four times the number of pixels compared to the resolution of the input video, which was 640×480 . This enhancement upgrades an NTSC grade camera to an HTDV

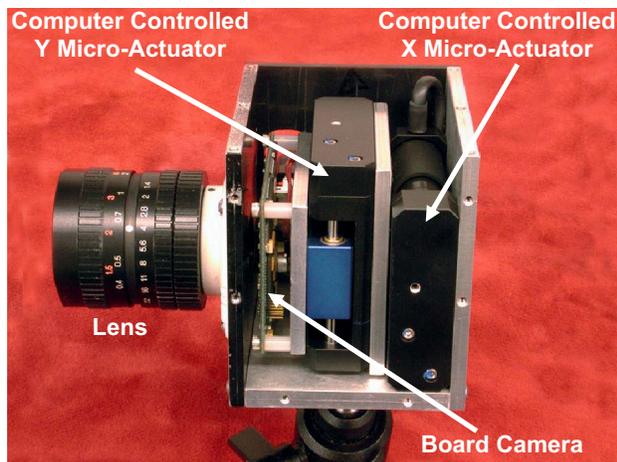


Figure 5: The jitter camera prototype shown with its cover open. The mechanical micro-actuators are used for shifting the board camera. The two actuators and the board camera are synchronized such that the camera is motionless during integration time.

grade camera while maintaining the depth of field and the frame-rate of the original camera.

With the recent advances in micro-electric mechanical systems (MEMS), it will hopefully be possible to embed the jitter mechanism within the detector chip, thus creating a jitter-detector.

5. Adaptive Super-resolution for Dynamic Scenes

Given a video sequence captured by a jitter camera, we would like to compute a high resolution video using super-resolution. We have chosen iterated-back-projection [10] as the super-resolution algorithm. Iterated-back-projection was shown in [5] to produce high quality results and is simple to implement for videos containing complex scenes. The main challenge in our implementation is handling multiple motions and occlusions. Failing to cope with these problems results in strong artifacts that render the output useless.

To address these problems, we compute the image motion in small blocks, and detect blocks suspected of having multiple motions. The adaptive super-resolution algorithm maximizes the use of the available data for each block.

5.1 Motion Estimation in the Presence of Aliasing

The estimation of image motion should be robust to outliers, which are mainly caused by occlusions and multiple motions within a block. To address this problem, we use the Tukey M-estimator error function [12]. The Tukey M-estimator depends on a scale parameter σ , the standard deviation of the gray-scale differences of correctly-aligned image regions (*inlier regions*).

Due to the under-sampling of the image, gray-scale image differences in the inlier regions are *dominated by aliasing*, and are especially significant near sharp image edges. Hence we approximate the standard deviation of the gray-scale differences

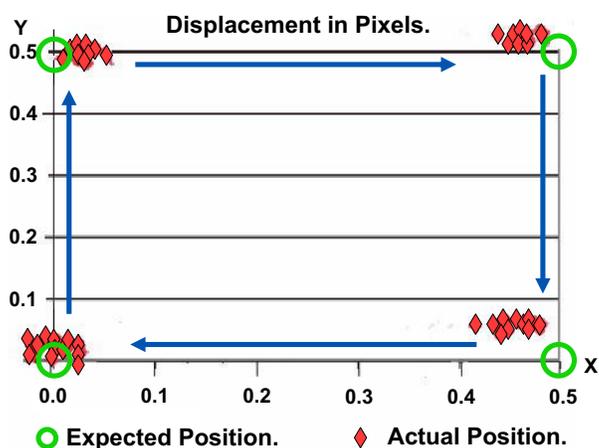


Figure 6: Accuracy of the jitter mechanism. The detector moves one step at a time along the path shown by the blue arrows. The green circles show the expected position of exactly half a pixel displacement and the red diamonds show the actual position over multiple cycles. We can see that the accuracy was less than tenth of a pixel. We can also see that the jitter mechanism returns very accurately to its zero position, hence prevents excessive error accumulation over multiple cycles.

σ in each block from the standard deviation of the aliasing σ_a in the block, as $\sigma = \sqrt{2}\sigma_a$. This approximation neglects the influence of noise, and makes the simplifying assumption that the aliasing effects in two aligned blocks are statistically uncorrelated. In the following we describe the approximation for the standard deviation of the aliasing in each block σ_a , using results on the statistics of natural images.

Let f be a high resolution image, blurred and decimated to obtain a low resolution image g :

$$g = (f * h) \downarrow$$

where $*$ denotes convolution and \downarrow denotes subsampling. Let s be a perfect rect low pass filter. The aliasing in g is given by:

$$(f * h - f * s * h) \downarrow = f * h * (\delta - s) \downarrow$$

The band-pass filter $h * (\delta - s)$ can hence be used to simulate aliasing. For the motion estimation, we need to estimate σ_a , the standard deviation of the response of this filter to blocks of the *unknown* high resolution image. We use the response of this filter to the aliased low resolution input images to estimate σ_a . Let σ_0 be the standard deviation of the filter response to an input block. Testing with a large number of images, we found that σ_a can be approximated to be a linear function of σ_0 . Similar results for non-aliased images were shown by Simoncelli [20] for various band-pass filters at different scales. For blocks of size 16×16 pixels the linear coefficient was in the range $[0.5, 0.7]$. In the experiments, we set $\sigma_a = 0.7\sigma_0$ which was sufficient for our purpose.

5.2 Adaptive Data Selection

We use the scale estimate σ from the previous section to differentiate between blocks with a single motion and blocks that



Input video frame



Blocks usage map

Figure 7: Adaptation of the super-resolution algorithm to moving objects and occlusions. The image on top shows one frame from a video sequence of a dynamic scene. The image on bottom is a visualization of the number of valid blocks, from four frames, used by the algorithm in each block. We darkened blocks where the algorithm used less than four valid blocks due to occlusions.

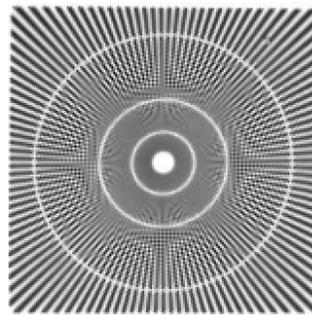
may have multiple motions and occlusions. A block in which the SSD error exceeds 3σ is excluded from the super-resolution calculation. In order to double the resolution (both horizontally and vertically) three additional valid blocks are needed for each block in the current frame. Depending on the timing of the occlusions, these additional blocks could be found in previous frames only, in successive frames only, both, or not at all. We therefore search for valid blocks in both temporal directions and select the blocks which are valid and closest in time to the current frame.

In blocks containing a complex motion, it may happen that less than four valid blocks are found within the temporal search window. In this case, although the super-resolution image is under-constrained, iterated-back-projection produces reasonable results [5]. Figure 7 shows an example from an outdoor video sequence containing multiple moving objects. On bottom is a visualization of the number of valid blocks used for each block in this frame. Blocks where less than four valid blocks were used are darkened.

6. Experiments

We tested resolution enhancement with our jitter camera for both static and dynamic scenes. The input images were obtained

Raw video from jitter camera



Super-resolution output

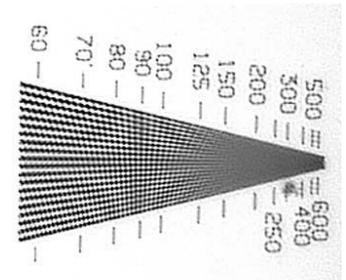
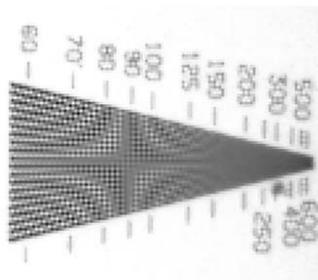
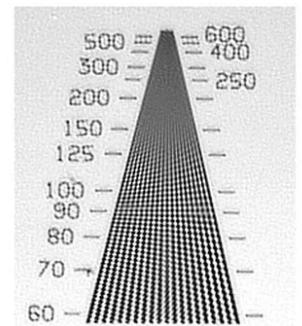
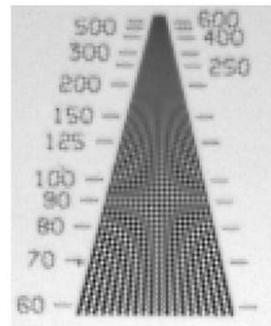
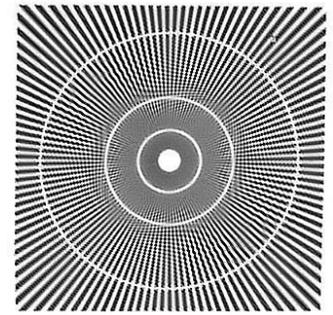


Figure 8: Resolution test using a standard Kodak test target. The left column shows angular, vertical and horizontal resolution test targets that were captured by the jitter camera (one of four input images). The right column shows the super-resolution results. Note the strong aliasing in the input images and the clear separation between lines in the super-resolution result images.

from the raw Bayer-pattern samples using the proprietary de-mosaicing algorithm provided by the camera manufacturer [16]. The images were then transformed to the *CIE-Lab* color space, and the super resolution algorithm [10] was applied to the L-channel only. The low resolution (a,b)-chroma channels were linearly interpolated and combined with the high resolution L-channel.

6.1 Resolution Tests

The resolution enhancement was evaluated quantitatively using a standard Kodak test target. The input to the super-resolution algorithm was four frames from a jitter-camera video sequence.

Figure 8 shows angular, vertical and horizontal test patterns. the aliasing effects are clearly seen in the input images, where the line separation is not clear even at the lower resolution of 60 lines per inch. In the computed super-resolution images the spatial resolution is clearly enhanced in all angles and it is possible to resolve separate lines well above 100 lines per inch.

6.2 Dynamic Video Tests

Several experiments were conducted to test the system's performance in the presence of moving objects and occlusions

Figure 9 shows magnified parts of an outdoor and an indoor sequences. Static objects, such as the crossing pedestrians sign in the first row, and the boat in the third row are significantly enhanced, revealing new details. The second row shows a particular dynamic scene with many moving objects. One can see that the adaptive super-resolution algorithm has increased the resolution of stationary objects while preserving or increasing the resolution of moving objects.

7. Conclusions and Future Work

Super-resolution algorithms can improve spatial resolution. However, their performance depends on various factors in the camera imaging process. We showed that motion blur causes significant degradation of super-resolution results, even when the motion blur function is known. The proposed solution is the jitter camera, a video camera capable of sampling the space-time volume without introducing motion blur. Applying a super-resolution algorithm to jitter camera video sequences significantly enhances their resolution.

As image detectors are becoming smaller and lighter and thus require very little force to jitter and with recent advances it may be possible to manufacture jitter cameras with the jitter mechanism embedded inside the detector chip. Jittering can then be added to regular video cameras as an option that enables a significant increase of spatial resolution while keeping other factors such as frame-rate unchanged.

Motion blur is only one factor in the imaging process. By considering other factors, novel methods for sampling the space-time volume can be developed, resulting in further improvements in video resolution. In this paper, for example, we limited the detector to a regular sampling lattice and to regular temporal sampling. One interesting direction can be the use of different lattices and different temporal samplings. We therefore consider the jitter camera to be a first step towards a family of novel camera designs that better sample the space-time volume to improve not only spatial resolution, but also temporal resolution and spectral resolution.

References

[1] S. Baker and T. Kanade. Limits on super-resolution and how to break them. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 24(9):1167–1183, September 2002.

[2] B. Bascle, A. Blake, and A. Zisserman. Motion deblurring and super-resolution from an image sequence. In *European Conf. on Computer Vision*, pages II:573–582, 1996.

[3] M. Ben-Ezra, A. Zomet, and S.K. Nayar. Jitter camera: High resolution video from a low resolution detector. Technical Report CUCS-016-04, Department of Computer Science, Columbia University, April 2004.

[4] J.R. Bergen, P. Anandan, K.J. Hanna, and R. Hingorani. Hierarchical model-based motion estimation. In *European Conf. on Computer Vision*, pages 237–252, 1992.

[5] D. Capel and A. Zisserman. Super-resolution enhancement of text image sequences. In *Int. Conf. Pattern Recognition*, pages Vol I: 600–605, September 2000.

[6] M.C. Chiang and T.E. Boult. Efficient super-resolution via image warping. *Image and Vision Computing*, 18(10):761–771, July 2000.

[7] Pixera Corporation. Director, <http://www.pixera.com/>.

[8] M. Elad and A. Feuer. Restoration of a single superresolution image from several blurred, noisy, and undersampled measured images. *IEEE Trans. Image Processing*, 6(12):1646–1658, December 1997.

[9] Physik Instrumente. M-111 micro translation stage, <http://www.physikinstrumente.de/>.

[10] M. Irani and S. Peleg. Improving resolution by image registration. *GMIP*, 53:231–239, 1991.

[11] Z. Lin and H.Y. Shum. Fundamental limits of reconstruction-based superresolution algorithms under local translation. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 26(1):83–97, January 2004.

[12] P. Meer, D. Mintz, D.Y. Kim, and A. Rosenfeld. Robust regression methods for computer vision: A review. *International Journal of Computer Vision*, 6(1):59–70, 1991.

[13] Minolta. Dimage-a1, <http://www.dpreview.com/reviews/minoltadimagea1/>.

[14] A.J. Patti, M.I. Sezan, and A.M. Tekalp. Superresolution video reconstruction with arbitrary sampling lattices and nonzero aperture time. *IEEE Trans. Image Processing*, 6(8):1064–1076, August 1997.

[15] A. Rav-Acha and S. Peleg. Restoration of multiple images with motion blur in different directions. In *IEEE Workshop on Applications of Computer Vision*, pages 22–28, 2000.

[16] Point Grey Research. Dragonfly camera, <http://www.ptgrey.com/>.

[17] R.R. Schultz and R.L. Stevenson. Extraction of high-resolution frames from video sequences. *IEEE Trans. Image Processing*, 5(6):996–1011, June 1996.

[18] E. Shechtman, Y. Caspi, and M. Irani. Increasing space-time resolution in video. In *European Conf. on Computer Vision*, page I: 753 ff., 2002.

[19] H. Shekarforoush and R. Chellappa. Data-driven multichannel superresolution with application to video sequences. *Journal of the Optical Society of America*, 16(3):481–492, March 1999.

[20] E.P. Simoncelli. Modeling the joint statistics of images in the wavelet domain. *SPIE*, 3813:188–195, July 1999.

Raw video from jitter camera

Super-resolution output



Figure 9: Jitter camera super-resolution for scenes with dynamic and stationary objects. The left column shows the raw video input from the jitter camera and the right column shows the super-resolution results. The first row shows a part of a scene with mostly stationary objects. Note the significant resolution enhancement of the pedestrians on the sign, and the fine texture of the tree branches. The second row shows a particular dynamic scene with many moving objects. Note the enhancement of the face of the walking women (center) and the kid on the scooter (left). The third row shows a scene with a large stationary object (boat) and a large moving object (woman's head). As expected, the resolution enhancement is better for the boat.