

A Hybrid Camera for Motion Deblurring and Depth Map Super-Resolution

Feng Li Jingyi Yu

Department of Computer and Information Sciences
University of Delaware
Newark, DE 19716, USA
{feli, yu}@cis.udel.edu

Jinxiang Chai

Department of Computer Science
Texas A&M University
College Station, TX 77843, USA
jchai@cs.tamu.edu

Abstract

We present a hybrid camera that combines the advantages of a high resolution camera and a high speed camera. Our hybrid camera consists of a pair of low-resolution high-speed (LRHS) cameras and a single high-resolution low-speed (HRLS) camera. The LRHS cameras are able to capture fast-motion with little motion blur. They also form a stereo pair and provide a low-resolution depth map. The HRLS camera provides a high spatial resolution but also introduces severe motion blur when capturing fast moving objects. We develop efficient algorithms to simultaneously motion-deblur the HRLS image and reconstruct a high resolution depth map. Our method estimates the motion flow in the LRHS pair and then warps the flow field to the HRLS camera to estimate the point spread function (PSF). We then deblur the HRLS image and use the resulting image to enhance the low-resolution depth map using joint bilateral filters. We demonstrate the hybrid camera in depth map super-resolution and motion deblurring with spatially varying kernels. Experiments show that our framework is robust and highly effective.

1. Introduction

With the recent advances in digital imaging, the use of high resolution, high speed, or high dynamic range cameras has become a common practice. However, thus far no single image sensor can satisfy the diverse requirements of all industrial camera applications today. For example, high-speed cameras can capture fast motion with little motion blur, but require expensive sensing, bandwidth and storage. The image resolution in these cameras is often much lower than many commercial still cameras. This is mainly because the image resolution needs to linearly scale with the exposure time [2] to maintain the Signal-to-Noise Ratio (SNR), i.e., higher speed maps to a lower resolution. In addition, the relatively low bandwidth on usual interfaces like

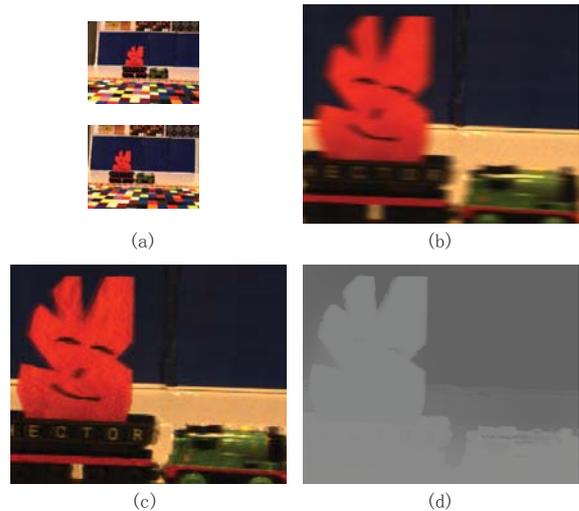


Figure 1. Motion deblurring and depth map super-resolution results using our hybrid camera. (a) shows two frames from the LRHS cameras, (b) shows a cropped region of the HRLS image. (c) showed the motion deblurred result. (d) shows the reconstructed high resolution depth map.

USB 2.0 or FireWire IEEE 1394a also restricts the image resolution especially when streaming videos at 100 to 200 frame/second.

In this paper, we present a hybrid sensor that combines the advantages of a high resolution camera and a high speed camera. Our imaging system consists of a pair of low-resolution high-speed (LRHS) cameras and a single high-resolution low-speed (HRLS) camera. The LRHS cameras can capture fast moving objects with very little motion blur. They also form a stereo pair and can generate a coarse depth map. The HRLS camera, on the other hand, provides a high spatial resolution yet is vulnerable to motion blur. We develop efficient algorithms to simultaneously motion-deblur the HRLS image and reconstruct a high resolution depth map.

Our method first recovers a low resolution depth map

from the LRHS pair and combines it with each LRHS’s motion flow to estimate the point spread function (PSF) in the HRLS image. Unlike most previous motion deblurring methods that only support spatially-invariant kernels, our method identifies the regions with different motion flow and individually deblurs each region using the Richardson-Lucy (R-L) algorithm. Once we deblur the HRLS image, we then use it to upsample the low resolution depth map using joint bilateral filters. An example of the results generated by our hybrid camera is shown in Figure 1.

2. Related work

Our work is mostly related to depth super-resolution methods and image deblurring algorithms.

Depth Map Super-resolution The problem of depth map super-resolution is relatively new in computer vision. Diebel and Thrun [7] modeled the disparity as a Markov Random Field on the low resolution depth map and then used a high resolution image to enhance the map using conjugate gradient algorithm. G. Piella [18] proposed a region-based approach that combines aspects of both object and pixel-level fusion for image and depth map super-resolution. Sawhney et al. [20] proposed a hybrid stereo camera that uses image-based rendering methods to accurately warp the high resolution image to the low resolution one. Kopf et al. used a high resolution image as a prior to upsample the low resolution exposure map and depth map using joint bilateral filters [12].

Recently, Yang et al. [24] proposed to combine the low resolution depth map obtained from a range sensor and a high resolution image to create a super-resolution depth map. They assume no parallax between the range sensor and the camera and their scene objects are generally static. The cameras in our system, in contrast, have a relatively large baseline. We also aim to recover the high resolution depth map of fast moving objects.

Image Deblurring The problem of image-deblurring has been well studied in the image processing community. Most previous methods have focused on recovering the blur kernel, or the point spread function (PSF). The classical Wiener filter [25] attempts to estimate the PSF from the image statistics and uses regularization to compute the inverse kernel. Computer vision methods such as graph cuts [19] and belief propagation [21] have also been used to recover nearly optimal deblurred images. Tull and Katsaggelos [23] proposed an iterative restoration approach to further improve the quality of deblurred images. Jia et al. [10] improved a short exposure noisy image by using color constraints observed in a long exposure photo without solving for the PSF. Sun et al. [26] used a blurred/noisy image pair to more accurately estimate the blur kernel. Most single-image deblurring algorithms assume the PSF is spatially-invariant. An exception is the work by Levin [13] in which

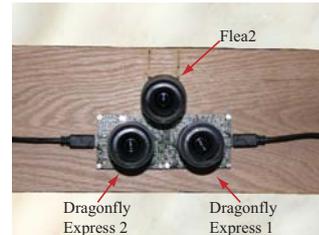


Figure 2. The prototype of our proposed hybrid camera. It consists of a rig of three cameras: a PointGrey Flea2 that serves as the HRLS camera and two PointGrey Dragonfly Express that serve as the LRHS cameras.

the image is segmented into several layers with different kernels, each assumed to have a constant motion velocity.

Hardware solutions have also been proposed to reduce the motion blur. These include the techniques based on lens stabilization and sensor stabilization. Adaptive optical elements controlled by inertial sensors have been used to compensate for camera motion [6, 16]. Liu and Gamal [14] proposed to capture multiple images on a CMOS sensor to construct a high dynamic range and motion deblurred image.

Closest to our design, Ben-Ezra and Nayar [2] proposed a hybrid imaging system that consists of a LRHS and a HRLS camera. They assume that the motion blurs are caused by the shaking of the camera and they used the LRHS to estimate the PSF in the HRLS. It is possible to apply their system to deblurring moving objects [3]. However, their method only supports spatially invariant kernels (e.g., the scenes of constant depth or with a single moving object) whereas our system is able to deblur images with spatially varying kernels such as the scenes that consist of multiple objects moving along different trajectories and/or at different speeds (Figure 6). In addition, we use the deblurred HRLS image to enhance the low resolution depth map.

3. System Setup and Algorithm Overview

Our prototype hybrid camera uses a pair of PointGrey Dragonfly Express cameras as the LRHS cameras. The Dragonfly Express captures images of resolution 320x240 at 120 fps. We also position a PointGrey Flea2 camera on top of the two Dragonfly cameras. The Flea2 serves as the HRLS camera and captures images of resolution 1024x768 at 7.5 fps. We use the software solution provided by PointGrey to synchronize the LRHS cameras and the HRLS camera so that every HRLS frame synchronizes with 16 LRHS frames. These three cameras are connected to two SIIG FireWire 800 3-Port PCIe cards. We attach all three camera modules to a wood plate and mount it on 2 tripods, as shown in Figure 2. Each camera uses a micro-lens with 4.8mm fo-

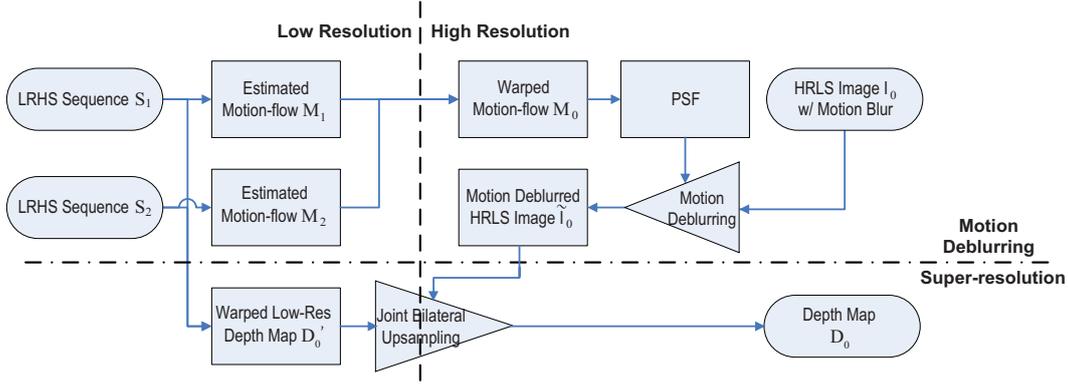


Figure 3. The processing pipeline using our hybrid camera for motion deblurring and depth map super-resolution. We first estimate the motion flows M_1 and M_2 in two LRHS sequences. We then warp the flow to the HRLS camera as M_0 . To motion deblur the HRLS image, we estimate the PSF from M_0 and use the Richardson-Lucy algorithm to deblur the image. To construct a high resolution depth map D_0 , we warp the low resolution depth map estimated from the LRHS cameras to the downsampled HRLS and use joint bilateral filters to upsample the D'_0 .

cal length.

An overview of our motion deblurring and depth map super-resolution framework is shown in Figure 3. We assume each frame I_0 in the HRLS camera maps to two LRHS sequence $S_1(t)$ and $S_2(t)$ ($t = 1 \dots K$ and $K = 16$ in our case). We first estimate the motion flow fields M_1 and M_2 in S_1 and S_2 and warp them onto I_0 as M_0 . Next, we estimate the PSF in I_0 from M_0 and apply the R-L algorithm to deblur I_0 . Recall that I_0 corresponds to K consecutive LRHS frames, therefore, I_0 can be deblurred using K different PSFs, each derived from the relative motions with respect to a specific LRHS frame. We use $\tilde{I}_0(t)$ to represent the deblurred result of I_0 for frame t . To generate the super-resolution depth map, we first compute a low resolution depth map $D_L(t)$ between $S_1(t)$ and $S_2(t)$. We then warp $D_L(t)$ to the downsampled HRLS camera as $D'_0(t)$ and then upsample $D'_0(t)$ using a joint bilateral filter, whose spatial kernel is defined on $D'_0(t)$ and range kernel defined on $\tilde{I}_0(t)$.

4. Motion Deblurring

In this section, we show how to use our hybrid camera for efficient motion deblurring.

4.1. Estimating Motion Flow

We first partition each frame in the LRHS camera into multiple foreground regions Ω_i^f and a background region Ω^b . To do so, we simply take a background image without any foreground objects and then use fore/background subtraction followed by graph-cut [5] to extract the foreground regions. We then group all foreground pixels into separate connected regions. We assume that each region has homogeneous motion and the regions do not overlap. We also use

the estimated disparity map (Section 5.1) to establish correspondences between the foreground regions in two LRHS cameras. This allows us to individually motion-deblur each foreground region using the method described in Section 4.3. Notice that since our system captures a sequence of images, it is also possible to directly composite the background image by applying a median filter across the frames.

To estimate the motion flow in each region Ω_i^f , we assume an affine motion model between two consecutive frames although we can concatenate several successive frames to form more complex motions. We apply a multi-resolution iterative algorithm [4] to minimize the energy function:

$$\arg \min_p \sum_x [I(w(x; p)) - I'(x)]^2 \quad (1)$$

$$w(x; p) = \begin{pmatrix} p_1x + p_2y + p_3 \\ p_4x + p_5y + p_6 \end{pmatrix}$$

where p corresponds to the motion flow parameter, w is the warping function. In our case, we estimate the motion flow in each LRHS camera, i.e., $I = S_j(t)$ and $I' = S_j(t+1)$, $j = 1, 2$. At each frame, we use the Gauss-Newton method to find the optimal motion [1].

4.2. Motion Warping

Once we estimate the motion flow for every foreground region Ω_i^f in each LRHS camera, we warp the the motion flow to the HRLS camera. Denote $M_1(t)$ and $M_2(t)$ as the estimated motion flow sample in S_1 and S_2 at frame t . To compute the motion flow sample $M_0(t)$ in the HRLS image I_0 , we first establish correspondences between $S_1(t)$ and $S_2(t)$. We use SIFT feature detection [15] to process $S_1(t)$ and $S_2(t)$ and then perform a global matching. To remove

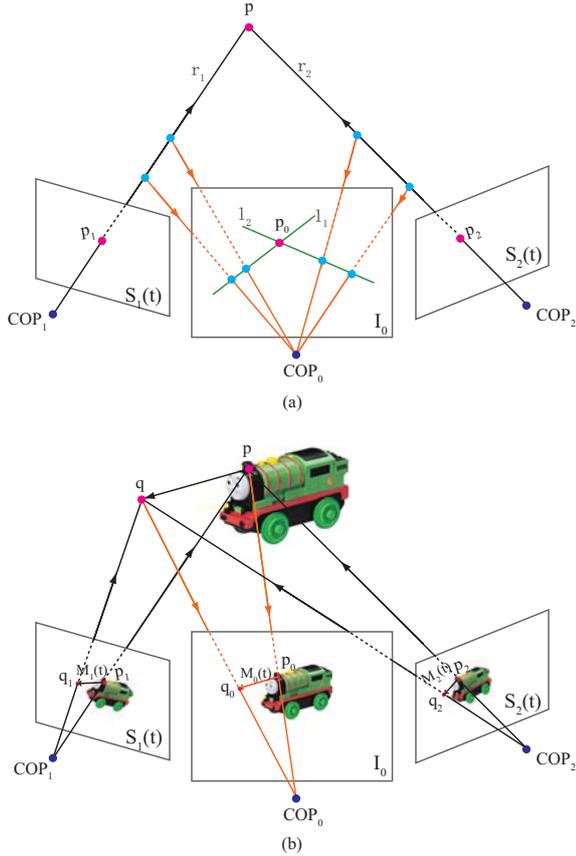


Figure 4. Motion estimation. (a) For each correspondence pair p_1 and p_2 in the LRHS pair, we project the ray that passes through p_1 in $S_1(t)$ onto I_0 as l_1 . We similarly project p_2 as l_2 in I_0 . Finally, we compute the intersection point of l_1 and l_2 to find the corresponding point p_0 . (b) To estimate the motion flow of p_0 in I_0 , we use the motion flow in the LRHS cameras to find q_0 and treat p_0q_0 as its motion flow.

outliers, our algorithm uses RANSAC with projective transformations as its precondition.

Given a pair of corresponding feature points p_1 and p_2 in $S_1(t)$ and $S_2(t)$, we connect p_1 with S_1 camera’s center-of-project (COP) to form a ray r_1 . We then project r_1 onto the HRLS image I_0 as a line l_1 . We apply a similar process to p_2 to obtain line l_2 . Finally we intersect l_1 and l_2 to obtain the p_0 , as is shown in Figure 4(a).

To estimate the motion flow of p_0 in the HRLS camera I_0 , we assume p_0 ’s corresponding point p_1 in $S_1(t)$ moves to q_1 by motion flow $M_1(t)$ and p_2 in $S_2(t)$ moves to q_2 by $M_2(t)$. Similar to the way we find p_0 , we then combine q_1 and q_2 to find their corresponding point q_0 in I_0 . We use the displacement between p_0 and q_0 as a motion flow sample. To improve robustness, we average multiple motion flow samples to compute $M_0(t)$.

Since our method combines the correspondences and the

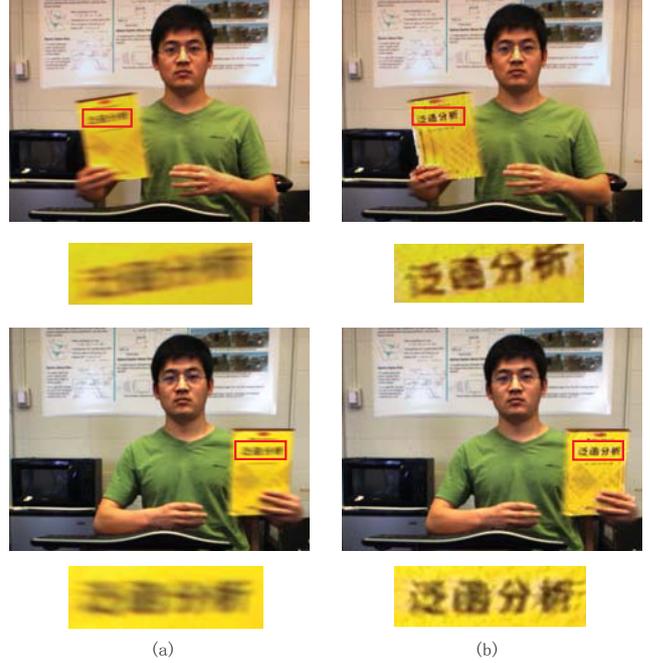


Figure 5. Motion deblurring results in a dynamic scene. The yellow book was quickly passed over from one hand to the other. (a) shows two original HRLS frames and (b) shows the deblurred results using our approach.

LRHS’s motion flow, our camera is able to model more complex motions than the system proposed by Ben-Ezra and Nayar. In [2], the motion flow in the HRLS camera is assumed to be identical to the one estimated from the LRHS camera. This is only valid for planar motions. Our method separately estimates the motion flow for each foreground region Ω_i^f and warps it using the correspondences. Therefore, we can compute the motion flow for each foreground objects in I_0 , even if they lie on different depth or move along different trajectories. Notice that the accuracy of our motion warping relies heavily on camera calibrations. In our implementations, we use the method in [27] to calibrate all 3 cameras.

4.3. PSF estimation and Image deconvolution

Recall that every HRLS frame maps to K ($K = 16$) LRHS frames, therefore, we can compute the PSF with respect to each frame t . To do so, we concatenate all K discrete motion samples $M_0(t)$, $t = 1 \dots K$ to form a single motion path. Ben-Ezra and Nayar proposed to use the motion centroid to construct a Voronoi tessellation for smoothly interpolating the motion path. Notice that our LRHS camera captures at a very high frame rate, therefore, each motion sample usually only covers 3 to 5 pixels in the HRLS image. To accurately estimate the kernel, we first appropriately align the motion path with the center of the kernel,

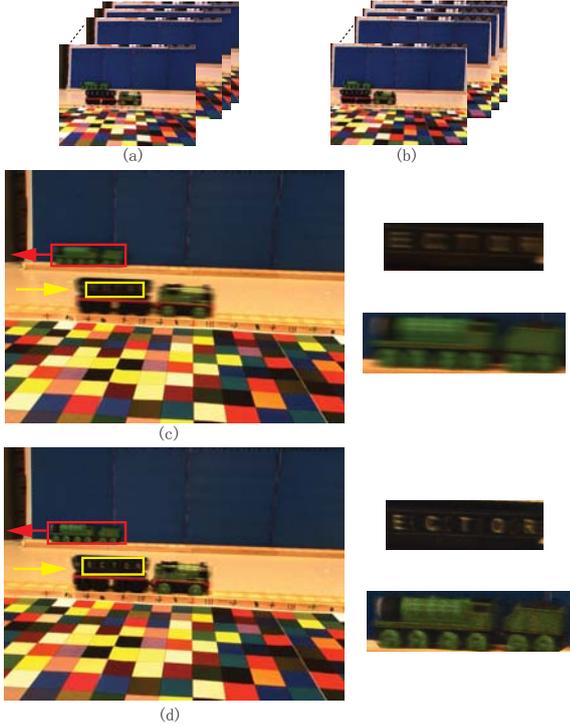


Figure 6. Motion deblurring results with spatially varying kernels. (a) and (b) show the two LRHS sequences. (c) shows the HRLS image with strong motion blur. (d) shows the motion deblurred results using our method. Notice the front and back train are moving in opposite directions at different speeds. Our method is able to motion deblur both objects.

then resample each $M_0(t)$ into N subpixels ($N = 50$ in our experiment), we further count, for each pixel p covered by $M_0(t)$ in the kernel, the number of subpixels N_p falling into p . Finally, we use this count N_p to estimate the weight of entry p in the kernel and normalize the kernel as the PSF.

Once we estimate the PSF, we can deblur the HRLS image using existing image deconvolution algorithms [9, 25]. In our implementation, we choose to use the Richardson-Lucy (R-L) iterative deconvolution algorithm. The R-L deconvolution always produces non-negative gray level values and works better than linear methods if the blur kernel is known and the noise level in the image is low. Before applying the R-L algorithm, we first estimate the masks for each of the foreground regions in the HRLS image that are obtained by using the technique presented in section 4.1 and 4.2. These masks are used to composite deblurred results into a common background image. In Figure 5 and 6, we show the motion results on captured dynamic scenes.

5. Depth Map Super-resolution

Next, we show how to use the deblurred HRLS image for generating super-resolution depth maps.

5.1. Initial Depth Estimation

We first use the two LRHS cameras and apply the graph-cut algorithm [11] to estimate a low-resolution disparity map with respect to the left LRHS camera. We then convert it to the depth map D_L . Next, we warp D_L to the downsampled HRLS camera as D'_0 . We adopt a similar notation as in [27]: every pixel p_1 in D_L is mapped to its corresponding pixel p_0 in D'_0 as:

$$p_0 = (\mathbf{A}_0 \mathbf{R}_0) (\mathbf{A}_1 \mathbf{R}_1)^{-1} (p_1 - \mathbf{A}_1 \mathbf{T}_1) + \mathbf{A}_0 \mathbf{T}_0 \quad (2)$$

where p is homogeneous coordinate of the pixel as $(sx, sy, s)^T$, s represents the depth of the point, \mathbf{A}_1 and \mathbf{A}_0 are the camera intrinsic matrices of the LRHS camera S_1 and the downsampled HRLS camera. \mathbf{R} and \mathbf{T} are the extrinsic to the cameras. Here we downsample the HRLS to have the same spatial resolution as LRHS camera to reduce the missing data (holes) in the reprojected depth map.

A sample depth map D_L is shown in Figure 7(b). Notice the depth map is incomplete because of the large baseline between the two LRHS cameras. This is less problematic since we focus on capturing the depth map of the moving foreground objects. However, the foreground can exhibit some serious artifacts. For example, in Figure 8, the left boundary of the red foreground object in the depth map partially merges with the background and the silhouettes of the toy train lack fine details. In Section 5.2, we show how to use the joint bilateral upsampling to reduce these artifacts.

We choose to warp the depth map to the HRLS camera instead of warping the HRLS image to the LRHS camera mainly because the depth-based warping inevitably corrupts the quality of the image by introducing holes near occlusion boundaries. Since our goal is to construct a high resolution depth map, it is more preferable to maintain the quality of the high resolution image than the low resolution depth map.

5.2. Joint Bilateral Upsampling

To enhance the warped depth map, we use the recently proposed joint bilateral filters. The basic bilateral filter is an edge-preserving filter [22] that uses both a spatial filter kernel and a range filter kernel evaluated on the data values themselves. A joint bilateral filter chooses the range filter from a second guidance image. It can effectively combine flash/no-flash pairs [17, 8], upsample the low resolution exposure maps, and enhance the low resolution range maps [12, 24]. In this paper, we also use the joint bilateral filter for depth map super-resolution.

Our joint bilateral filter combines the low resolution depth map D'_0 and the motion deblurred high resolution image \tilde{I}_0 . It computes the value at each pixel p in the high

resolution depth map D_0 as:

$$D_0(p) = \frac{1}{W_p} \sum_{q \in \Theta \setminus \Gamma} G_s(\|p - q\|) G_r(|\tilde{I}_0(p) - \tilde{I}_0(q)|) D'_0(q) \quad (3)$$

$$W_p = \sum_{q \in \Theta \setminus \Gamma} G_s(\|p - q\|) G_r(|\tilde{I}_0(p) - \tilde{I}_0(q)|)$$

where G_s is the spatial kernel centered over p and G_r is the range kernel centered at the image value at p in \tilde{I}_0 ; Θ is the spatial support of the kernel G_s . W_p is the normalization factor. Since the warped low resolution depth map D'_0 contains holes, we exclude the points Γ that correspond to the holes.

To emphasize on the color difference in the range kernel, we choose

$$|\tilde{I}_0(p) - \tilde{I}_0(q)| = \max(|r_p - r_q|, |g_p - g_q|, |b_p - b_q|) \quad (4)$$

where r, g, b are the color channel of the pixel in \tilde{I}_0 . In Figure 7(c), we show the reprojected low resolution depth map of resolution 320x240 and the enhanced high resolution depth map in Figure 7(d). The joint bilateral filter significantly improves the quality of the depth map by filling in the missing holes and partially correcting the inaccurate boundaries (e.g., the background between the engine and carriage of the toy train in Figure 7). However, the resulting boundaries still appear blurry due to the spatial Gaussian. Although we can re-apply the joint bilateral filters on the upsampled depth image to further improve the boundary, we implement a simple scheme based on the observation that with a large enough spatial support Θ , there should be enough pixels that have the similar color and depth to pixel p . To do so, we apply a joint median filter: we first pick N pixels from Θ that have the closest color to p and then apply the median filter on the depth values of these N pixels. In our experiment, we choose N to be 20% of all pixels in Θ .

6. Results and Discussion

Figure 2 shows the prototype of our hybrid camera. The LRHS camera captures images of resolution 320x240 at 120 fps and HRLS camera captures images of resolution 1024x768 at 7.5 fps. The two LRHS are positioned about 7.5cm away from each other and the HRLS camera is placed above the LRHS cameras.

In Figure 5 and 6, we show the motion deblurring results using our hybrid camera. In Figure 5(a), we apply motion deblurring algorithm to a dynamic scene: the yellow book was quickly passed over from the right hand to the left hand. Notice, the texts on the book are unrecognizable in the original HRLS frames, but are clearly readable in the deblurred results shown in Figure 5(b). The hand holding the textbook is also effectively deblurred. In Figure 6, we

use the hybrid camera to deblur spatially varying kernels. The foreground and the background toy trains are moving in opposite directions at different speeds. Since our method separately estimates the PSF for each individual train using the method discussed in Section 4, we are able to effectively deblur both objects.

In Figure 7, we validate our joint bilateral filter method for depth map super resolution in a static scene. We position a toy train on top of the rails. Notice the toy train has complex contours. If we simply upsample the low resolution depth map obtained from the LRHS cameras, the details of the contours disappear. Using a joint bilateral upsampling, we are able to recover these fine details and partially correct the erroneous boundaries in the depth map (e.g., the V shaped contour between the engine and the carriage).

Finally, we show using the hybrid camera to simultaneously motion deblur the HRLS image and reconstruct the high resolution depth map. Figure 1 shows the results using our system in a toy train scene. We mount a star shaped object on top of the toy train to emphasize the accuracy of the recovered depth maps. The LRHS and HRLS inputs are shown in Figure 1(a) and Figure 8(a), and the deblurred result is shown in Figure 1(c) and Figure 8(b). Our method effectively reduces the motion blur by recovering the text on the toy train and the contour of the star-shaped object, although we also observe some ringing artifacts caused by the R-L deconvolution.

In Figure 8, we show the depth super resolution results using joint bilateral filters. We choose a large enough kernel size σ_s for the spatial Gaussian to fill in the missing data. In this example, we set σ_s to be 35. We normalize the depth map to have intensity between [0 1] and we set σ_r of the range Gaussian to be 0.02. Figure 8(c) shows the depth map from the LRHS pair and Figure 8(d) shows the warped result. The reprojected depth map is very noisy and contains many holes. Using the joint bilateral filter, the depth map is significantly enhanced (Figure 8(e)). The seams at the occlusion boundaries are further enhanced (Figure 8(f)) using the joint median filter as described in Section 5.2. Close-ups are shown at the bottom of Figure 8. The typical up-sampling rate in our experiments is 9. Due to the camera calibration requirement, we did not further downsample the LRHS camera to demonstrate a higher up-sampling rate, although it can easily be achieved if we use a different HRLS camera with a higher resolution.

7. Conclusions and Future Work

We have presented a hybrid camera that combines the advantages of a high resolution camera and a high speed camera. Our hybrid camera consists of a pair of low-resolution high-speed (LRHS) cameras and a single high-resolution low-speed (HRLS) camera. The LRHS cameras are able to capture fast motion with little motion blur. They also form

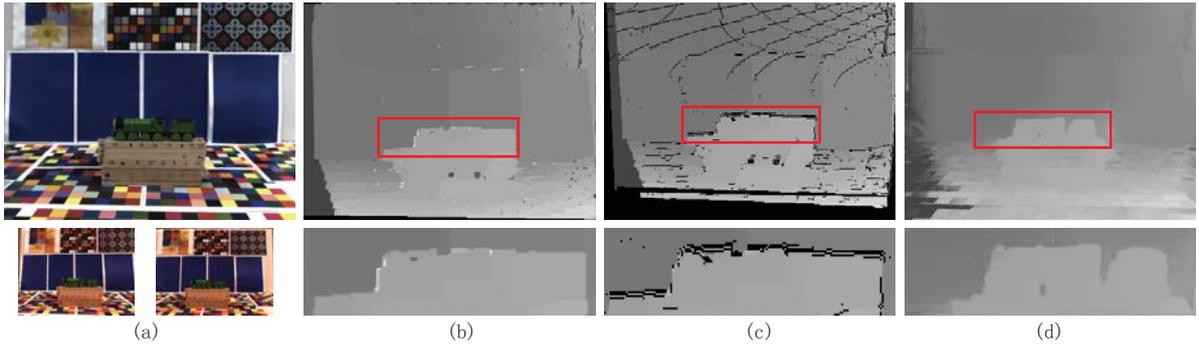


Figure 7. Depth map super-resolution results for a static scene. (a) shows the input images from the LRHS and HRLS cameras. (b) shows the depth map from the LRHS camera pair. The result is upsampled using bilinear interpolation for comparison. (c) shows the reprojected low resolution depth map onto the HRLS camera. The warped map is very noisy and contains holes. (d) shows the upsampled results using the joint bilateral filters.

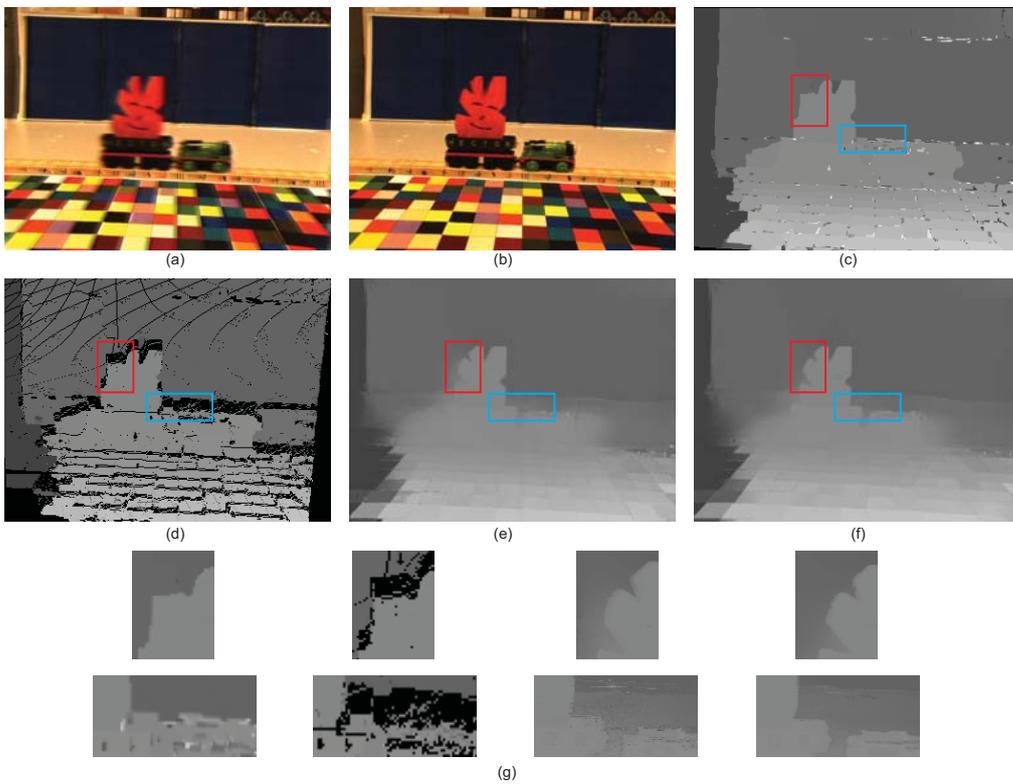


Figure 8. Motion deblurring and depth map super-resolution using the hybrid camera. (a) shows the motion blurred image from the HRLS camera. (b) shows the motion deblurred result using the method described in Section 4. (c) shows the low resolution depth map. We upsample the map using bilinear interpolation to compare it with the bilateral filtered results. (d) is the reprojected depth map. (e) shows the upsampling result using joint bilateral filters. (f) shows further improved results using the joint median filter described in Section 5.2. (g) shows closeup images of the depth maps.

a stereo pair and provide a low-resolution depth map. The HRLS camera provides a high spatial resolution but also introduces severe motion blur when capturing fast moving objects.

We have developed efficient algorithms to simultane-

ously motion-deblur the HRLS image and reconstruct a high resolution depth map. Our method estimates the motion flow in the LRHS pair and then warps the flow field to the HRLS camera to estimate the PSF. We then deblurred the HRLS image and used the resulting image to enhance

the low-resolution depth map using joint bilateral filters. We have demonstrated the hybrid camera's usefulness in depth map super-resolution and motion deblurring with spatially varying kernels.

Our future work includes experimenting the hybrid camera for outdoor scenes, using our system to generate videos with extremely high spatial and temporal resolutions, and varying the exposure time in the two LRHS cameras for producing high resolution and high dynamic range videos.

Acknowledgement

The authors thank Dr. Sylvain Paris for his invaluable inputs and Dr. Ben-Ezra for his useful suggestions. This work has been supported by the National Science Foundation under grant NSF-MSPA-MCS-0625931.

References

- [1] S. Baker and I. Matthews. Lucas-kanade 20 years on: A unifying framework. *Int. J. Comput. Vision*, 56(3):221–255, 2004. 3
- [2] M. Ben-Ezra and S. Nayar. Motion deblurring using hybrid imaging. *Computer Vision and Pattern Recognition, 2003. Proceedings. 2003 IEEE Computer Society Conference on*, 1:I–657–I–664 vol.1, 18-20 June 2003. 1, 2, 4
- [3] M. Ben-Ezra and S. Nayar. Motion-based motion deblurring. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 26(6):689–698, June 2004. 2
- [4] J. R. Bergen, P. Anandan, K. J. Hanna, and R. Hingorani. Hierarchical model-based motion estimation. *ECCV '92: Proceedings of the Second European Conference on Computer Vision*, pages 237–252, 1992. 3
- [5] Y. Boykov and G. Funka-Lea. Graph cuts and efficient n-d image segmentation. *Int. J. Comput. Vision*, 70(2):109–131, 2006. 3
- [6] Canon. What is optical image stabilizer? <http://www.canon.com/bctv/faq/optis.html>, 2006. 2
- [7] J. Diebel and S. Thrun. An application of markov random fields to range sensing. *NIPS*, 2005. 2
- [8] E. Eisemann and F. Durand. Flash photography enhancement via intrinsic relighting. *ACM Trans. Graph.*, 23(3):673–678, 2004. 5
- [9] R. Fergus, B. Singh, A. Hertzmann, S. T. Roweis, and W. Freeman. Removing camera shake from a single photograph. *ACM Trans. Graph.*, 25:787–794, 2006. 5
- [10] J. Jia, J. Sun, C. Tang, and H. Shum. Bayesian correction of image intensity with spatial consideration. *ECCV*, 3:342–354, 2004. 2
- [11] V. Kolmogorov and R. Zabih. Multi-camera scene reconstruction via graph cuts. *European Conference on Computer Vision*, May 2002. 5
- [12] J. Kopf, M. F. Cohen, D. Lischinski, and M. Uyttendaele. Joint bilateral upsampling. *SIGGRAPH '07: ACM SIGGRAPH 2007 papers*, page 96, 2007. 2, 5
- [13] A. Levin. Blind motion deblurring using image statistics. *NIPS*, 2006. 2
- [14] X. Liu and A. E. Gamal. Simultaneous image formation and motion blur restoration via multiple capture. *ICASSP '01: Proceedings of the Acoustics, Speech, and Signal Processing, 2001. on IEEE International Conference*, pages 1841–1844, 2001. 2
- [15] D. Lowe. Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vision*, 20:91–110, 2003. 3
- [16] Nikon. Precise camera-shake compensation at every angle. http://www.nikon.co.jp/main/eng/portfolio/about/technology/nikon_technology/vr_e/index.htm, 2005. 2
- [17] G. Petschnigg, R. Szeliski, M. Agrawala, M. Cohen, H. Hoppe, and K. Toyama. Digital photography with flash and no-flash image pairs. *ACM Trans. Graph.*, pages 664–672, 2004. 5
- [18] G. Piella. A general framework for multiresolution image fusion: from pixels to regions. *Information Fusion*, 4(4):259–280, 2003. 2
- [19] A. Raj and R. Zabih. A graph cut algorithm for generalized image deconvolution. *ICCV '05: Proceedings of the Tenth IEEE International Conference on Computer Vision*, pages 1048–1054, 2005. 2
- [20] H. S. Sawhney, Y. Guo, K. Hanna, R. Kumar, S. Adkins, and S. Zhou. Hybrid stereo camera: an ibr approach for synthesis of very high resolution stereoscopic image sequences. *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 451–460, 2001. 2
- [21] M. Tappen, B.C. Russell, and W. Freeman. Efficient graphical models for processing images. *Computer Vision and Pattern Recognition*, 2:II–673–II–680 Vol.2, 2004. 2
- [22] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. *ICCV '98: Proceedings of the Sixth International Conference on Computer Vision*, page 839, 1998. 5
- [23] D. Tull and A. Katsaggelos. Iterative restoration of fast-moving objects in dynamic image sequences. *Optical Engineering*, 35(12):3460–3469, December 1996. 2
- [24] Q. Yang, R. Yang, J. Davis, and D. Nister. Spatial-depth super resolution for range images. *Computer Vision and Pattern Recognition, 2007. CVPR '07. IEEE Conference on*, pages 1–8, 17-22 June 2007. 2, 5
- [25] Y. Yitzhaky, I. Mor, A. Lantzman, and N. S. Kopeika. Direct method for restoration of motion-blurred images. *Journal of the Optical Society of America A: Optics, Image Science, and Vision*, 15(6):1512–1519, 1998. 2, 5
- [26] L. Yuan, J. Sun, L. Quan, and H.-Y. Shum. Image deblurring with blurred/noisy image pairs. *SIGGRAPH '07: ACM SIGGRAPH 2007 papers*, page 1, 2007. 2
- [27] Z. Zhang. A flexible new technique for camera calibration. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22(11):1330–1334, 2000. 4, 5