

Stereo Video Deblurring

Anita Sellent^{1,2}(✉), Carsten Rother¹, and Stefan Roth²

¹ Technische Universität Dresden, Dresden, Germany

anita.sellent@tu-dresden.de

² Technische Universität Darmstadt, Darmstadt, Germany

Abstract. Videos acquired in low-light conditions often exhibit motion blur, which depends on the motion of the objects relative to the camera. This is not only visually unpleasing, but can hamper further processing. With this paper we are the first to show how the availability of stereo video can aid the challenging video deblurring task. We leverage 3D scene flow, which can be estimated robustly even under adverse conditions. We go beyond simply determining the object motion in two ways: First, we show how a piecewise rigid 3D scene flow representation allows to induce accurate blur kernels via local homographies. Second, we exploit the estimated motion boundaries of the 3D scene flow to mitigate ringing artifacts using an iterative weighting scheme. Being aware of 3D object motion, our approach can deal robustly with an arbitrary number of independently moving objects. We demonstrate its benefit over state-of-the-art video deblurring using quantitative and qualitative experiments on rendered scenes and real videos.

Keywords: Object motion blur · Scene flow · Spatially-variant deblurring

1 Introduction

Stereo is one of the oldest areas of computer vision research [1]. Interestingly, the arrival of mass-produced active depth sensors [2] seems to have renewed interest also in passive stereo systems. In contrast to active depth sensors, stereo cameras are also applicable in outdoor environments. Due to their more general applicability, stereo cameras are gaining increased adoption, for example in autonomous driving [3]. Remarkably, the availability of stereo image pairs also helps in the estimation of temporal correspondences: On the KITTI optical flow benchmark [4], the best performing algorithms [5,6] are indeed scene flow algorithms that jointly estimate depth and 3D motion from stereo videos. Part of their advantage stems from an increased robustness to adverse imaging conditions [6]. One such adverse imaging condition is a shortage of light. In low-light conditions,

Electronic supplementary material The online version of this chapter (doi:[10.1007/978-3-319-46475-6_35](https://doi.org/10.1007/978-3-319-46475-6_35)) contains supplementary material, which is available to authorized users.

the exposure time often needs to be increased to obtain a reasonable signal-to-noise-ratio. But when either the camera, or the objects in the scenes are moving during exposure time, this results in motion blurred images.

Motion blur is not only unsatisfactory to look at, it can also disturb further image-based processing, *e.g.* in tasks such as panorama stitching [8] or barcode recognition [9]. In stereo video setups, viewpoint-dependent motion blur hinders a post-capture adjustment of the baseline, the acquisition and visualization of 3D point clouds (see Fig. 1 for an example) or the control of tele-operated robots in the presence of rapid robot and/or object motion.




Fig. 1. Application of stereo video deblurring: Given 2 consecutive stereo frames (a), our deblurring approach allows to estimate sharp textures from stereo video input with motion blur. Rendering scene flow geometry with the blurred input image as a colored point-cloud from a new point of view produces an unnatural motion blur (b). Our stereo video deblurring algorithm can remove the blur (c).

In this paper we address the challenge of deblurring stereo videos. In contrast to the substantial literature on removing camera shake [10–15], we aim to deal with the more general case of camera *and* object motion. In case of independent motions, mixed pixels at motion boundaries yield significant complications. Removing such spatially-variant blur is extremely challenging when attempted from single images [16, 17], but video input helps to significantly increase robustness [7, 18]. Unlike previous work, we leverage stereo video to obtain substantially improved and more robust deblurring results. In our approach, we exploit 3D scene flow in various ways and *make the following contributions*: (i) We show that 3D scene flow can improve video deblurring by providing more accurate motion estimates. In particular, we exploit piecewise rigid scene flow [6], which yields an over-segmentation of the image into planar patches that move with a rigid 3D motion (Figs. 2b and c). (ii) We demonstrate that the resulting piecewise homographies allow to directly induce blur matrices. Thereby, we take into account that the projection of a rigid 3D motion yields non-linear motion trajectories in 2D (Fig. 3, Table 1). We find that this leads to superior deblurring results compared to inducing the blur matrices from an optical flow field [7] (Figs. 2d to f). (iii) We apply the homography-induced blur matrices in a robust deblurring procedure that attenuates the effects of motion discontinuities using

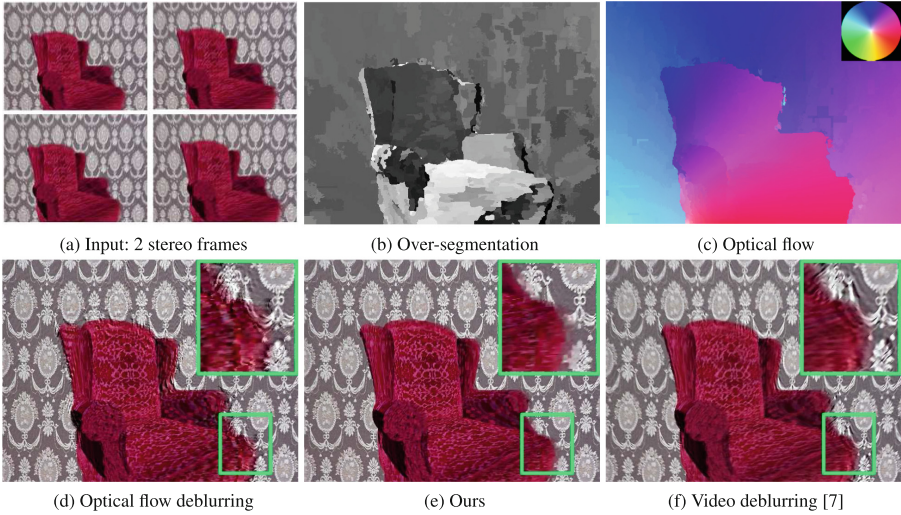


Fig. 2. Stereo video deblurring: For two consecutive frames of a synthetic stereo video (a) we use the scene flow approach of Vogel *et al.* [6] to compute an over-segmentation into planar patches with constant 3D rigid body motion (b). Projecting the 3D motion onto the image plane yields optical flow (c), which our baseline algorithm uses to deblur a reference frame (d). Exploiting the homographies from the 3D motion and object boundary information from the over-segmentation, our full approach obtains sharp images avoiding ringing and boundary artifacts (e). Our result is also clearly sharper than state-of-the-art video deblurring [7] (f)




Fig. 3. Descriptiveness of homography-based blur kernels: Using 3D rigid body motion to generate blur kernels, we can faithfully express, *e.g.*, yaw motion (a), while kernels constructed with spatially varying 2D displacement vectors fields [7] only yield an approximation (b). Approximation errors (c) are also present close to the rotation axis where motions are small (extremely large yaw angle and all intensities scaled for better visibility)

an iterative weighting scheme; Initial motion discontinuities are obtained from 3D scene flow. We demonstrate the superiority of the proposed stereo video deblurring over state-of-the-art monocular video deblurring using experiments on synthetic data as well as on real videos.

Table 1. Overview of the different sources of motion information used for video deblurring: When pure 2D correspondence is considered (top two rows), the induced blur kernels are only approximate, as motion trajectories are assumed to be linear. Exploiting homographies from scene flow allows us to capture the fact that rigid 3D object motion leads to non-linear trajectories

Motion information	# of frames	# of views	Linear approximation of trajectories
Optical flow	2	1	Yes
2D projection of scene flow	2	2	Yes
Homographies	2	2	No

2 Related Work

The goal of this work is to obtain sharp images from stereo videos containing 3D camera and object motion. Of course, in principle blind deblurring could be applied to each frame individually. However, blind motion deblurring from a single image is a highly underconstrained problem, as blur parameters and sharp image have to be estimated from a single measurement. To cope with spatially-variant blur due to the 3D motion of the camera, single image deblurring approaches frequently use homographies [19–21]. In contrast, we apply homographies to describe spatially-variant object motion blur. Single image object motion deblurring approaches keep the number of parameters manageable by either choosing the motion of a region from a very restricted set of spatially-invariant box filters [22, 23], assuming it to have a spatially-invariant, non-parametric kernel of limited size [16], or to be representable by a discrete set of basis kernels [24]. Approaches that rely on learning spatially-variant blur are also limited to a discretized set of detectable motions [17, 25]. Kim *et al.* [26] consider continuously varying box filters for every pixel, but rely heavily on regularization.

Connecting deblurring and depth estimation, Xu and Jia [27] successfully apply stereo correspondence estimation to motion-blurred stereo frames to support blind image deblurring. Lee and Lee [28], Arun *et al.* [29], and Hu *et al.* [30] estimate sharp images and depth jointly. However, all these approaches assume the scene to be static and camera motion to be the only source of motion blur.

Cho *et al.* [31] deblur images of independently moving objects. The multiple input images of their algorithm are unordered, and a piecewise affine registration between the images, as well as the motion underlying the blur, has to be estimated. To restrict the parameter space, the blur kernels are assumed to be piecewise constant and linear.

Video deblurring approaches reduce the number of parameters through the assumption that the inter-frame and intra-frame motion are related by the duty cycle of the camera. He *et al.* [32] and Deng *et al.* [33] apply feature tracking of a single moving object to obtain 2D displacement-based blur kernels for deblurring. Wulff and Black [18] refine the latter approach and perform segmentation

into two layers, estimation of the affine motion parameters, as well as deblurring of each layer jointly. Relaxing the assumption of two layers and affine motion, Yamaguchi *et al.* [34] and Kim and Lee [7] employ optical flow to approximate spatially variant blur kernels for deblurring. Yamaguchi *et al.* [34] propose deblurring based on the flow estimates from the blurry images. Kim and Lee [7] iteratively refine flow estimation and deblurred video frames by minimizing a joint energy. The latter method represents the state-of-the-art in video deblurring and is used for comparison in the experimental section. To the best of our knowledge, exploiting stereo video for deblurring has not been considered in the literature before.

Correspondence estimation on stereo video sequences can be improved by estimating stereo correspondences and optical flow jointly as 3D scene flow [35–37]. In our approach we build on the piecewise rigid scene flow by Vogel *et al.* [6] for the following reasons. First, it provides us with explicit 3D rotations and translations that we employ for accurate blur kernel construction. Second, through over-segmentation into planar patches, it also delivers occlusion information, which we use as initialization for our boundary-aware object motion deblurring. A general problem in object motion deblurring is that object boundaries with mixed foreground and background pixels can lead to severe ringing artifacts (see Fig. 2). Explicit segmentation and α -matting [18, 38] can prevent this effect, but requires restrictive assumptions on the number of moving objects. To handle general scenes with an arbitrary number of objects, we extend the robust outlier handling of Chen *et al.* [39] to spatially-variant deblurring based on scene flow, and apply it to the mixed pixels at object boundaries.

In contrast to the aforementioned deblurring approaches, Cho *et al.* [40] deblur hand-held video under the assumption that patches are sharp in some frames of the video. However, in the case of autonomous robots or objects passing the field of view with high speed, this assumption does not hold. Joshi *et al.* [41] attach additional inertial measurement units to the camera, but this does not account for object motion. An additional low-resolution, high frame-rate camera can provide complex motion kernels [38], but does not provide depth estimates in the way a stereo camera can.

3 Blurred Image Formation in Stereo Video

Inducing Blur Matrices from 3D Rigid Object Motions. Due to the finite exposure time τ of our stereo video camera, each frame of each camera is blurred. Our goal is to find a sharp image I_{t_0} for a reference camera at time t_0 . We base our approach on the scene flow of Vogel *et al.* [6], and likewise assume that the scene can be approximated with planar patches that undergo a 3D rigid body motion. If an object in the scene is non-planar, this assumption leads to an over-segmentation of the object into spatially adjacent patches (see Fig. 2b). Considering video frames where the exposure time is naturally limited by the frame rate, we additionally assume that the motion of each patch is constant during the exposure time of two consecutive frames. Note that a constant rigid

motion in 3D does not necessarily imply that its 2D projection is constant; the projection may, *e.g.* in the case of a rotation, be constantly accelerated. However, our assumption excludes rapidly changing motions such as vibrations.

Constant 3D rigid body motion can be expressed as a homogeneous 4×4 matrix

$$M = \begin{pmatrix} R & T \\ \mathbf{0} & 1 \end{pmatrix} \quad (1)$$

with a rotation matrix $R \in \mathbb{R}^{3 \times 3}$ and a translation vector $T \in \mathbb{R}^3$. To enable our highly accurate blur kernel description, we rewrite $M = \exp(\theta\xi)$ as matrix exponential, where $\theta \in \mathbb{R}$ describes the rotation angle and ξ is a 4×4 matrix that is determined by the rotation axis and the translation, see [42, 43]. With M describing the motion between time instants t_0 and t_1 , the constant 3D motion between two arbitrary time instants t_a and t_b is given as

$$M_{t_b, t_a} = \exp\left(\frac{t_b - t_a}{t_1 - t_0} \theta \xi\right). \quad (2)$$

In a piecewise planar scene approximation, the 3D planes of the patches at time t are defined via their scaled normals n_t . All points P on the plane satisfy the equation $P^T n_t = 1$, where P^T is the transposed of P . We can relate a moving 3D point to its corresponding pixel location on the image plane via the camera geometry. Given the calibration matrix K of the reference camera and its location T_K , the projection from a 3D plane to the image plane at time t can be written in homogeneous coordinates as $Pr_t = K - KT_K n_t^T$, see *e.g.* [6].

Under the assumption of color constancy, two sharp images of the reference camera (with hypothetical infinitesimal exposure) at different times are connected via

$$I_{t_a}(x) = I_{t_b}({}^{t_b}H^{t_a}x) \quad \text{where} \quad {}^{t_b}H^{t_a} = Pr_{t_b} M_{t_b, t_a} Pr_{t_a}^{-1}. \quad (3)$$

With this notation, a blurry image pixel x in the interior of a patch is formed from the reference image as

$$\hat{B}(x) = \int_{t_0 - \frac{\tau}{2}}^{t_0 + \frac{\tau}{2}} I_t(x) dt = \int_{t_0 - \frac{\tau}{2}}^{t_0 + \frac{\tau}{2}} I_{t_0}({}^{t_0}H^t x) dt, \quad (4)$$

where

$${}^{t_0}H^t = Pr_{t_0} \exp(-t\theta\xi) Pr_t^{-1} \quad (5)$$

is a homography that can be computed exactly from camera geometry, normal, and motion. To put it differently, a 3D point that is projected to x on the image plane describes a certain trajectory on the image plane during the exposure time. If the 3D point follows a rigid body motion, the homography $({}^{t_0}H^t)^{-1}$ allows us to exactly describe this 2D trajectory. In contrast, optical flow based methods [7, 24, 44], employ 2D optical flow vectors to generate I_t via forward warping. Thus the trajectory of a point on the image plane is approximated by a 2D line that is traversed with constant velocity. As optical flow is spatially

variant, the trajectories may change for each pixel, hence induce blur kernels with a curved shape. However, more complex motions such as rotations can only be *approximated*, Fig. 3. In our approach, the description of trajectories due to 3D rigid body motions is *exact*. As our experiments show, this also results in more faithfully deblurred images.

By discretizing the integration over time with $\delta t = \frac{\tau}{N}$ (we fix $N = 70$) and using bilinear image interpolation, we can obtain a discretized version of Eq. (4) for vectorized reference images as $\hat{B}(x) = A_x \mathbf{I}_{t_0}$. Here, A_x denotes a sparse row vector that depends on the homography estimated at pixel x . Stacking the blur vectors A_x for each pixel, we obtain our homography-based blur matrix A leading to $\hat{\mathbf{B}} = A \mathbf{I}_{t_0}$.

Motion Boundaries. If only scene points from the same plane contribute to the color $B(x)$ of the measured blurred image at point x , the image formation model of Eq. (4) is exact. If at time t a scene point with a different motion contributes to $B(x)$ we should also use the corresponding homography. However, within an object, the planar patches are adjacent in space and move consistently. Therefore, we approximate the blur with the row vector A_x induced by the homography of x at t_0 . At motion boundaries, the homographies are very different and as pixels of foreground and background mix, transparency effects occur. While such effects can be modeled, taking them into account requires precise localization of the motion boundaries, which is very challenging. Instead, we exclude motion boundaries from the deblurring process by means of an iterative approach. In each iteration, we downweight pixels with a high difference between image formation model and measured image and try to find a sharp image that explains the remaining pixels. Under the assumption of additive Gaussian noise, we use the residual to compute a weight for each pixel as

$$w_n(x) = \exp\left(-k_\sigma \|B(x) - A_x \mathbf{I}_{t_0}^{n-1}\|^2\right), \quad (6)$$

where $\mathbf{I}_{t_0}^{n-1}$ denotes the current estimate of the sharp (color) image. For normalized images we set $k_\sigma = 4000/3$ as default value. In the first iteration we initialize w_0 with the binary occlusion information from the scene flow. As Fig. 4 shows, the weights converge quickly. Some pixels in the image that were initially suppressed as motion boundaries are included in deblurring at a later iteration. More importantly, other pixels where the image formation model is invalid are suppressed later on, which helps controlling ringing artifacts. Suppression may also happen due to some inaccuracies in the computed scene flow. In the experimental section, we will see how this property actually helps to improve deblurring results.

Deblurring. Theoretically, we could fill in the regions at motion boundaries during deblurring by using adjacent frames or information from the other camera. However, we found experimentally that correspondence estimation in these regions is too unreliable to produce visually pleasing results. Instead, we exploit that natural, sharp images follow a Laplacian distribution of their gradients [22].




Fig. 4. Downweighting of mixed pixels due to motion boundaries: Foreground and background mix at motion boundaries and violate our image formation model (a). At motion boundaries and locations of inaccurate flow estimates, the image formation model is downweighted to avoid ringing artifacts. We initialize these weights with the occlusion information provided by the scene flow (b) and refine them iteratively (c), (d)

In locations where the image formation model is unreliable, *e.g.*, at motion boundaries, we rely on this prior to provide the necessary regularization. Specifically, we obtain an estimate of the sharp reference frame by minimizing the energy

$$E(\mathbf{I}_{t_0}) = \sum_{x \in \Omega} \left\| w_n(x)(B(x) - A_x \mathbf{I}_{t_0}) \right\|^2 + \alpha \rho(\nabla \mathbf{I}_{t_0}(x)), \quad (7)$$

where $\Omega \subset \mathbb{N}^2$ is the image domain and the constant α is fixed to 0.001. Following prior work [22], we use the robust norm $\rho(c) = |c|^{0.8}$ for each color channel and gradient direction.

To solve the optimization problem in Eq. (7), we use iteratively reweighted least squares (IRLS) [45]. In each reweighting iteration, we compute the following weights

$$\rho_n(c) = \frac{1}{c} \frac{d\rho(c)}{dc} \approx \max(|c|, \epsilon)^{0.8-2} \quad \text{with} \quad \epsilon = 0.01 \quad (8)$$

for the smoothness term using the preceding image estimate $\nabla \mathbf{I}_{t_0}^{n-1}$. Then we minimize the least squares energy

$$E(\mathbf{I}_{t_0}, n) = \sum_{x \in \Omega} \left\| w_n(x)(B(x) - A_x \mathbf{I}_{t_0}^n) \right\|^2 + \alpha \|\rho_n \nabla \mathbf{I}_{t_0}^n(x)\|^2 \quad (9)$$

via conjugate gradients. We alternate between updating the occlusion weight w_n and the smoothness weight ρ_n . In all our experiments the weights converge quickly and only a few (≈ 10) iterations were needed in total.

To compute the 3D scene flow needed for our stereo video deblurring approach, we rely on the method of Vogel *et al.* [6]. The algorithm is originally designed for sharp images. However, its data term uses the census transform for comparing the warped images, which makes it quite robust to image blur. Of course, scene flow estimation will reach its limits for very strong motion blur. Experimentally, we find that by aggregating evidence in piecewise planar patches, the method yields a scene flow accuracy that turns out to work well in deblurring stereo videos of casual motion. As the following experiments will show, it

is crucial, however, to not only rely on the robust correspondence information, but to exploit the homographies to directly induce the blur kernels.

4 Experiments

To demonstrate the efficacy of the proposed stereo video deblurring, we perform experiments on synthetic images with known ground truth, as well as on real images. We capture the real video footage with a *Point Grey Bumblebee2* stereo color camera, which can acquire 640×480 images at a frame rate of 20 Hz. We use the internal calibration and supplied software to obtain rectified and demosaiced images. The exposure time of each image can be obtained from the camera software.

In all experiments, we compute scene flow using the publicly available implementation of [6]. We take the default parameters and scale them uniformly to account for the baseline difference between our stereo camera and the KITTI dataset [4] for which they were tuned. For the 640×480 image in Fig. 2 our approach requires 73 s to form the discretized blur matrix A . Using MATLAB to optimize Eq. (7) in 25 conjugate gradient steps and 10 IRLS iterations requires 69 s on an 8-core 4 GHz CPU.

4.1 Comparing Flow-Based Deblurring to Homography-Based Deblurring

We begin by applying the proposed stereo video deblurring to scenes without object discontinuities. In this way we can analyze the benefit of the homography-induced motion blur model in isolation. We create synthetic sequences by simulating various 3D motions (upward and forward translation, and a combination forward translation and yaw) of a planar, roughly fronto-parallel texture, see Fig. 5a.¹ A second test set consists of rigidly moving 3D objects rendered with a raytracer at very small time steps and averaged to give motion-blurred images (see Figs. 4a, 6a and 7a for the first image of the left view). We take the central




Fig. 5. Deblurring planar textures: For a planar texture blurred with 3D rigid body motion (a), deblurring with 2D spatially-variant ground truth displacements (b) yields ringing errors (c) that can be reduced by deblurring with our homography-based image formation model (d), (e)

¹ More textures and motions are evaluated in the supplemental material.

frame of each motion-blurred image as a sharp reference frame. For the rendered scenes motion discontinuities are known. In the first experiment, we disable the data term around any motion discontinuities by fixing the weights w_n in these areas to zero, see Fig. 6b for an example. As the image prior stays active, the boundaries are filled in smoothly as illustrated in Fig. 6d.

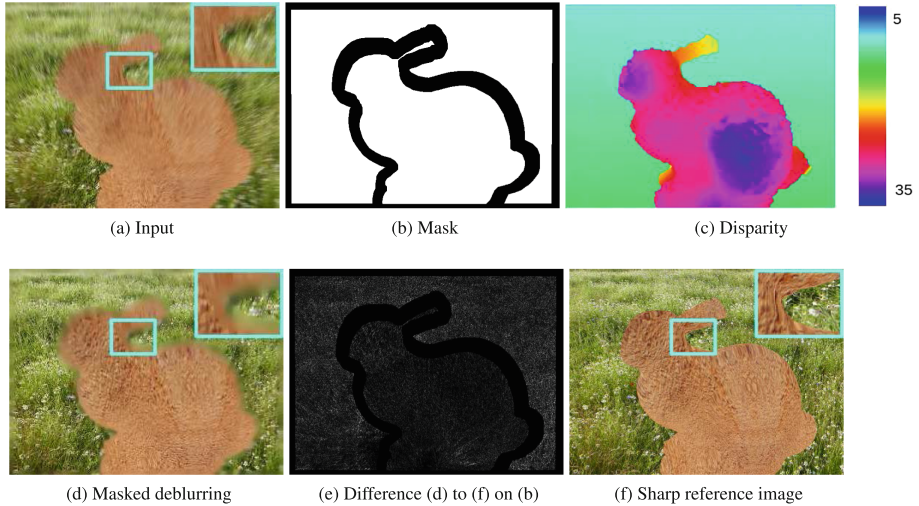


Fig. 6. Deblurring with masked discontinuities: Our raytraced stereo video frames contain independent object motion of non-planar objects (a). Through the estimated disparity we can assess the shape of the objects (c). Excluding given discontinuities (b) from the computation of the data term, invalid areas are filled in smoothly (d). The masked difference image (e) to the real sharp image (f) shows that homography-based deblurring has about the same error in planar as in curved surfaces, showing the effectiveness of the over-segmentation from the scene flow

We compare our homography-induced deblurring approach against deblurring with blur matrices generated from different 2D displacement fields. We use forward and backward 2D motion as described by Kim and Lee [7] and apply them in our IRLS deblurring framework. In particular, we use the known *ground truth 2D displacement*, the *2D initial optical flow* with which the scene flow is initialized [46] (*baseline deblurring*), and the *2D projection of the scene flow* to induce blur kernels. Table 1 summarizes these settings. Table 2 shows the peak-signal-to-noise-ratio (PSNR) of the deblurred images from the different methods. We observe that the PSNR of our homography-based stereo video deblurring outperforms the results of deblurring with ground truth 2D displacement in all cases of non-fronto-parallel motion. In these cases linear motion trajectories of constant velocity are an approximation. Blur matrices induced by homographies are more expressive and improve the results. Already, deblurring with the 2D projection of scene flow achieves a consistently higher PSNR than deblurring with the




Fig. 7. Raytraced scenes for evaluating object motion deblurring: The input images exhibit blur due to 3D object motion (a). Also when 3D homographies are used to induce blur kernels, mixed pixels at object boundaries cause some ringing artifacts (b). Iteratively downweighting the boundaries from the data term, our full stereo video deblurring (c) suppresses ringing and obtains considerably sharper images than state-of-the-art video deblurring (d). Please zoom in for detail

initial flow. Indeed, in the case of forward motion, also deblurring with the 2D projection of the scene flow outperforms deblurring with ground truth displacement. The estimated 2D displacement appears to be a better approximation to the linear, but accelerated trajectory of the 3D forward motion than the 2D ground truth displacement. Figures 5b and d show examples of deblurred images using the ground-truth 2D displacement and our homography-based approach. From the difference image between the results and the original sharp texture, Figs. 5c and e, we observe that the increase in PSNR is due to the mitigation of ringing effects throughout the image.

For the raytraced scenes the geometry of the moving objects is non-planar and the planarity assumption in our image formation model becomes an approximation. Figure 6 shows the estimated disparity of an object and the deblurred image obtained by masking out discontinuities. Looking at the difference image, Fig. 6e, we observe that the deblurring error for slightly curved surfaces is comparable to the performance on planar regions of the background, showing that the over-segmentation aids coping with curved surfaces.

Table 2. Deblurring without considering motion-discontinuity regions: For different motions of a planar texture (top) and moving 3D objects with masked object boundaries (bottom), we report the peak signal-to-noise ratio (PSNR) of the deblurred reference frame, the average endpoint error of the estimated motion (AEP), and the average disparity error (ADE) of the estimation. For all scenes the use of scene flow increases deblurring accuracy compared to using optical flow. For scenes with non-fronto-parallel motion (all except ‘*upward*’ and ‘*apples*’) homography-based object motion deblurring provides the best results (bold)

Blur kernel source	Ground truth 2D displacement	Initial optical flow [46]	2D projection of scene flow	3D homographies
	PSNR	AEP PSNR	AEP PSNR	ADE PSNR
upward	26.09	2.53 24.89	0.20 25.83	3.66 25.94
forward	34.12	0.09 34.45	0.10 34.55	0.19 34.74
forward + yaw	30.03	0.15 29.90	0.15 29.95	0.45 31.02
apples	34.36	0.48 29.39	0.81 34.74	4.95 33.33
bunny	25.01	0.62 24.87	0.54 25.09	0.50 25.19
chair	24.84	0.59 23.86	0.47 24.81	0.25 25.03
squares	25.97	2.32 22.69	0.97 25.82	0.65 27.21
triplane	27.43	1.16 27.11	0.49 27.27	0.12 28.30

For all rendered scenes where the disparity does not exhibit gross errors, we observe in Table 2 that 3D homography-based deblurring improves the PSNR clearly over any form of 2D deblurring. In the scene ‘*apples*’, Fig. 7a, 1st row, depth estimation fails with a mean disparity error of 4.95 pixels. In this situation the deblurring quality of homography-based deblurring drops below that of its 2D projection. Still, both outperform the results obtained with the initial optical flow. More importantly, as we will see below, the iterative weighting scheme for treating motion discontinuities can address such disparity estimation errors as well and lead to much improved results.

4.2 Full Algorithm with Motion Discontinuities

We now evaluate the performance of stereo video deblurring in the presence of object motion boundaries. We use the raytraced scenes from the previous experiment, but this time without providing ground-truth information on the motion discontinuities, Fig. 7a. Additionally, we use real images captured with a stereo camera attached to a motorized rail, Fig. 8a. The camera moves forward very slowly on the rail while we capture frames with maximal exposure time and frame rate. By averaging the frames, we obtain motion-blurred images. Comparison to the central frame of the averaged frame series allows for numerical evaluation. Finally, we capture scenes with arbitrarily moving objects for which only a visual evaluation is possible, Fig. 9a. As before we compare against 2D versions of our algorithm. Additionally, we compare against the state-of-the-art

video deblurring algorithm of Kim and Lee [7] that uses 3 consecutive monocular frames. We tuned their regularization parameter to obtain the most accurate results.

In Figs. 7b and c we first contrast homography-induced deblurring without and with handling of motion boundaries. When not taking into account motion boundaries explicitly, *i.e.* $w_n \equiv 1$, Fig. 7b, considerable ringing artifacts are the result, but they are successfully suppressed with our proposed iterative weighting scheme, Fig. 7c. This also becomes evident in the numerical evaluation when comparing the 3rd and 4th column of Table 3 (top)². For the real sequences in Fig. 8, boundary artifacts are generally less pronounced, as all objects in the scene are static and the camera moves toward the scene. However, as shown in Fig. 4, the discontinuity weight can still compensate errors in scene flow computation. One such example is the erroneous depth estimation in the ‘apples’ scene, which is disabled by the discontinuity weight. Similarly, also in the scenes with the motorized rail, our full object motion deblurring approach improves the PSNR compared to the basic homography approach, Table 3 (bottom).

When comparing to the state-of-the-art video deblurring method of Kim and Lee [7], we find that our stereo video deblurring approach yields significantly fewer ringing artifacts and considerably sharper results. This can be seen visually, comparing (d) to (c) of Figs. 7, 8 and 9, as well as quantitatively in Table 3. Interestingly, we find in Table 3 that IRLS deblurring with the 2D projection of the scene flow is already on par with video deblurring of Kim and Lee. 3D homography-based deblurring without boundary handling

Table 3. Deblurring with motion discontinuities: PSNR of deblurred synthetic scenes with motion discontinuities (top) and real scenes with the camera moving on a motorized rail (bottom). Our homography-based stereo video deblurring with motion boundary weighting (full) clearly outperforms monocular video deblurring with optical-flow induced blur kernels in all cases

	Initial optical flow [46]	2D projection of scene flow	3D homographies	Ours (full)	Kim and Lee [7]
apples	20.89	29.43	26.00	29.48	26.14
bunny	20.72	22.36	22.95	23.20	21.93
chair	19.03	21.57	22.19	23.36	21.78
squares	19.60	21.90	22.56	24.58	22.99
triplane	24.73	24.55	25.22	26.59	23.30
bottles	29.51	29.55	30.80	31.07	28.33
office	31.37	30.87	32.45	32.56	29.01
planar	30.71	31.21	32.33	32.82	30.01
toys	33.51	33.64	34.89	34.90	31.06

² In the supplemental material we also consider the Structural Similarity index [47] to compare the results.

improves on these result numerically already, highlighting the importance of our homography-induced blur kernels. Yet, our full homography-based object deblurring with motion boundary handling gives further numerical gains and a large visual improvement. Recall that the motion boundaries are initially obtained from the 3D scene flow, thus unique to our setting.




Fig. 8. Controlled camera motion for evaluating object motion deblurring: Our 3D deblurring (c) has less ringing artifacts than baseline deblurring with optical flow (b), and sharper results than video deblurring (d), in particular at the periphery of the images where motion is large

For the real scenes with independent object motion, Fig. 9, we observe that the optical flow-based approaches introduce ringing artifacts, particularly where strong gradients of the background coincide with the object boundary. Our stereo video deblurring algorithm can cope with this situation even in the presence of non-planar, non-rigidly moving objects such as the trousers (2nd row) are present.




Fig. 9. For real scenes with independent object motion (a), our novel stereo video deblurring approach (c) generates fewer ringing artifacts due to object boundaries than baseline deblurring with optical flow (b) and sharper images than video deblurring (d)

5 Conclusions and Future Work

We have proposed the first stereo video deblurring approach, which is based on an image formation model that exploits 3D scene flow computed from stereo video. For scenes with an arbitrary number of moving objects, we use an over-segmentation of the scene into planar patches to establish spatially-variant blur matrices based on local homographies. Our experiments on synthetic scenes and real videos show that deblurring with these homographies is more accurate than baseline methods based on 2D linear motion approximations, as well as the current state-of-the-art in video deblurring. Combined with our robust treatment of motion boundaries through an iterative weighting scheme, our approach obtains superior results also on real stereo videos with independently moving objects. In future work we would like to improve the performance of scene flow computation at motion boundaries such that we can benefit from another view to supply information near motion boundaries.

Acknowledgements. The research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007 - 2013)/ERC Grant Agreement No. 307942 and under the European Union’s Horizon 2020 research and innovation programme/ERC Grant Agreement No. 647769.

References

1. Longuet-Higgins, H.C.: A computer algorithm for reconstructing a scene from two projections. *Nature* **293**, 133–135 (1981)
2. Shotton, J., Girshick, R., Fitzgibbon, A., Sharp, T., Cook, M., Finocchio, M., Moore, R., Kohli, P., Criminisi, A., Kipman, A., Blake, A.: Efficient human pose estimation from single depth images. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**(12), 2821–2840 (2013)
3. Franke, U., Joos, A.: Real-time stereo vision for urban traffic scene understanding. In: *Intelligent Vehicles Symposium*, pp. 273–278 (2000)
4. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? The KITTI vision benchmark suite. In: *CVPR*, pp. 3354–3361 (2012)
5. Menze, M., Geiger, A.: Object scene flow for autonomous vehicles. In: *CVPR*, pp. 3061–3070 (2015)
6. Vogel, C., Schindler, K., Roth, S.: 3D scene flow estimation with a piecewise rigid scene model. *Int. J. Comput. Vis.* **115**(1), 1–28 (2015)
7. Kim, T.H., Lee, K.M.: Generalized video deblurring for dynamic scenes. In: *CVPR*, pp. 5426–5434 (2015)
8. Li, Y., Kang, S.B., Joshi, N., Seitz, S.M., Huttenlocher, D.P.: Generating sharp panoramas from motion-blurred videos. In: *CVPR*, pp. 2424–2431 (2010)
9. Yahyanejad, S., Strom, J.: Removing motion blur from barcode images. In: *CVPR*, pp. 41–46 (2010)
10. Fergus, R., Singh, B., Hertzmann, A., Roweis, S.T., Freeman, W.T.: Removing camera shake from a single photograph. In: *SIGGRAPH*, pp. 787–794 (2006)
11. Cho, S., Lee, S.: Fast motion deblurring. *ACM Trans. Graph.* **28**(5), 145:1–145:8 (2009)
12. Whyte, O., Sivic, J., Zisserman, A., Ponce, J.: Non-uniform deblurring for shaken images. In: *CVPR*, pp. 491–498 (2010)
13. Krishnan, D., Tay, T., Fergus, R.: Blind deconvolution using a normalized sparsity measure. In: *CVPR*, pp. 233–240 (2011)
14. Xu, L., Zheng, S., Jia, J.: Unnatural L_0 sparse representation for natural image deblurring. In: *CVPR*, pp. 1107–1114 (2013)
15. Michaeli, T., Irani, M.: Blind deblurring using internal patch recurrence. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) *ECCV 2014, Part III*. LNCS, vol. 8691, pp. 783–798. Springer, Heidelberg (2014)
16. Schelten, K., Roth, S.: Localized image blur removal through non-parametric kernel estimation. In: *ICPR*, pp. 702–707 (2014)
17. Couzinié-Devy, F., Sun, J., Alahari, K., Ponce, J.: Learning to estimate and remove non-uniform image blur. In: *CVPR*, pp. 1075–1082 (2013)
18. Wulff, J., Black, M.J.: Modeling blurred video with layers. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) *ECCV 2014, Part VI*. LNCS, vol. 8694, pp. 236–252. Springer, Heidelberg (2014)

19. Tai, Y.W., Tan, P., Brown, M.S.: Richardson-lucy deblurring for scenes under a projective motion path. *IEEE Trans. Pattern Anal. Mach. Intell.* **33**(8), 1603–1618 (2011)
20. Gupta, A., Joshi, N., Lawrence Zitnick, C., Cohen, M., Curless, B.: Single image deblurring using motion density functions. In: Daniilidis, K., Maragos, P., Paragios, N. (eds.) *ECCV 2010, Part I. LNCS*, vol. 6311, pp. 171–184. Springer, Heidelberg (2010)
21. Rajagopalan, A., Chellappa, R.: *Motion Deblurring: Algorithms and Systems*. Cambridge University Press, Cambridge (2014)
22. Levin, A.: Blind motion deblurring using image statistics. In: *NIPS*, pp. 841–848 (2006)
23. Chakrabarti, A., Zickler, T., Freeman, W.T.: Analyzing spatially-varying blur. In: *CVPR*, pp. 2512–2519 (2010)
24. Kim, T.H., Ahn, B., Lee, K.M.: Dynamic scene deblurring. In: *ICCV*, pp. 3160–3167 (2013)
25. Sun, J., Cao, W., Xu, Z., Ponce, J.: Learning a convolutional neural network for non-uniform motion blur removal. In: *CVPR*, pp. 769–777 (2015)
26. Kim, T.H., Lee, K.M.: Segmentation-free dynamic scene deblurring. In: *CVPR*, pp. 2766–2773 (2014)
27. Xu, L., Jia, J.: Depth-aware motion deblurring. In: *ICCP*, pp. 1–8 (2012)
28. Lee, H., Lee, K.: Dense 3D reconstruction from severely blurred images using a single moving camera. In: *CVPR*, pp. 273–280 (2013)
29. Arun, M., Rajagopalan, A., Seetharaman, G.: Multi-shot deblurring for 3D scenes. In: *CVPR*, pp. 19–27 (2015)
30. Hu, Z., Xu, L., Yang, M.H.: Joint depth estimation and camera shake removal from single blurry image. In: *CVPR*, pp. 2893–2900 (2014)
31. Cho, S., Matsushita, Y., Lee, S.: Removing non-uniform motion blur from images. In: *ICCV*, pp. 1–8 (2007)
32. He, X., Luo, T., Yuk, S., Chow, K., Wong, K.Y., Chung, R.: Motion estimation method for blurred videos and application of deblurring with spatially varying blur kernels. In: *ICCV*, pp. 355–359 (2010)
33. Deng, X., Shen, Y., Song, M., Tao, D., Bu, J., Chen, C.: Video-based non-uniform object motion blur estimation and deblurring. *Neurocomputing* **86**, 170–178 (2012)
34. Yamaguchi, T., Fukuda, H., Furukawa, R., Kawasaki, H., Sturm, P.: Video deblurring and super-resolution technique for multiple moving objects. In: Kimmel, R., Klette, R., Sugimoto, A. (eds.) *ACCV 2010, Part IV. LNCS*, vol. 6495, pp. 127–140. Springer, Heidelberg (2011)
35. Vedula, S., Baker, S., Rander, P., Collins, R., Kanade, T.: Three-dimensional scene flow. In: *ICCV*, vol. 2, pp. 722–729 (1999)
36. Quiroga, J., Devernay, F., Crowley, J.: Local/global scene flow estimation. In: *ICIP*, pp. 3850–3854 (2013)
37. Wedel, A., Cremers, D.: *Stereo Scene Flow for 3D Motion Analysis*. Springer Science & Business Media, London (2011)
38. Tai, Y.W., Du, H., Brown, M.S., Lin, S.: Correction of spatially varying image and video motion blur using a hybrid camera. *IEEE Trans. Pattern Anal. Mach. Intell.* **32**(6), 1012–1028 (2010)
39. Chen, J., Yuan, L., Tang, C.K., Quan, L.: Robust dual motion deblurring. In: *CVPR*, pp. 1–8 (2008)
40. Cho, S., Wang, J., Lee, S.: Video deblurring for hand-held cameras using patch-based synthesis. *ACM Trans. Graph.* **31**(4), 64:1–64:9 (2012)

41. Joshi, N., Kang, S.B., Zitnick, C.L., Szeliski, R.: Image deblurring using inertial measurement sensors. *ACM Trans. Graph.* **29**(4), 30:1–30:9 (2010)
42. Mei, C., Reid, I.: Modeling and generating complex motion blur for real-time tracking. In: *CVPR*, pp. 1–8 (2008)
43. Murray, R.M., Li, Z., Sastry, S.S.: *A Mathematical Introduction to Robotic Manipulation*. CRC Press, Boca Raton (1994)
44. Portz, T., Zhang, L., Jiang, H.: Optical flow in the presence of spatially-varying motion blur. In: *CVPR*, pp. 1752–1759 (2012)
45. Levin, A., Weiss, Y.: User assisted separation of reflections from a single image using a sparsity prior. *IEEE Trans. Pattern Anal. Mach. Intell.* **29**(9), 1647–1654 (2007)
46. Vogel, C., Roth, S., Schindler, K.: An evaluation of data costs for optical flow. In: Weickert, J., Hein, M., Schiele, B. (eds.) *GCPR 2013*. LNCS, vol. 8142, pp. 343–353. Springer, Heidelberg (2013)
47. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.* **13**(4), 600–612 (2004)