

Variational Optical Flow from Alternate Exposure Images

Anita Sellent¹, Martin Eisemann¹, Bastian Goldlücke², Thomas Pock³, Daniel Cremers², Marcus Magnor¹

¹Institut für Computergraphik, TU Braunschweig, Germany

²Department of Computer Science, University of Bonn, Germany

³Institute for Computer Graphics and Vision, Graz University of Technology, Austria

Email: {sellent, eisemann, magnor}@cg.tu-bs.de, {bastian, dcremers}@cs.uni-bonn.de,
pock@icg.tugraz.at

Abstract

Traditional optic flow algorithms rely on consecutive short-exposure images. In contrast, long-exposed images contain integrated motion information directly in form of motion blur. In this paper, we show how the additional information provided by a long exposure image can be used to improve robustness and accuracy of motion field estimation. Recently, an image formation model was introduced [23] that relates a long-exposure image to preceding and succeeding short-exposure images in terms of dense 2D motion and occlusion. We formulate the original two-step problem for motion and occlusion timings as a joint minimization problem and derive a global TV- L^1 energy functional that can be minimized efficiently and accurately. The approach is able to calculate highly accurate motion fields, assigning motion to occluded and disoccluded image regions in one joint optimization procedure.

1 Introduction

Estimating the dense motion field between two consecutive images is an old but still heavily investigated field of research. In order to solve the classical optical flow equation, the local time derivative needs to be numerically evaluated. Hence, most optical flow algorithms work best with pinpoint-sharp images as input, which depict a dynamic scene at two discrete points in time. If regarded individually, however, short exposed images capture no motion information.

From sampling theory, it is well known that this approach leads to temporal aliasing if maximum displacement exceeds one pixel [11]. The straightforward approach to avoid aliasing is to increase sampling rate, as has been done by calculating op-

tical flow from high-speed camera recordings [18]. If no high speed video equipment is available, adequate temporal pre-filtering is necessary. Since the motion field is a-priori unknown, instead, multi-scale optical flow methods pre-filter the images globally in the spatial domain [7]. This way, the entire image is low-pass filtered which is not the same as correct temporal filtering: high spatial frequencies should be suppressed only in those Fourier-domain regions where aliasing actually occurs, i.e., only in the direction of local motion.

Fortunately, there exists a simple way to achieve correct temporal pre-filtering: exposing images for an extended period of time. For moving objects, high frequencies in motion direction are suppressed in long exposure images. Reconstructing these high spatial frequencies in case of unknown motion is called blind image deblurring. Methods for blind image deblurring are probably as heavily investigated as those for motion detection [16].

Apart from circumventing the problem of temporal aliasing, long exposure images bear the advantage that occlusion enters into the image formation process. A scene point and its motion contribute to a motion-blurred image exactly for as long as the point is not occluded.

In [23] these observations were used to deal with optical flow estimation. Dense 2D motion fields were obtained by using three images with two distinct exposure times: the method requires images taken with alternating exposure intervals such that an intermediate, long-exposed image is enframed by two short-exposure images, Fig. 1. Already the straightforward pointwise L^2 optimization scheme of [23] is able to calculate promising motion fields from these informations.

In this paper, we build on the idea of [23] and extend it to achieve more robust results. The main




Figure 1: Alternately exposed images: (a) exposure timing diagram of (b) a short exposure image I_1 followed by (c) a long-exposed image I_B and (d) another short-exposed image I_2 .

contribution of this paper is to embed the alternate exposure image formation model in a global TV- L^1 energy formulation framework and to estimate motion and occlusion time jointly. The L^1 norm of the data-term ensures more robustness of the minimization to noise in the intensity values of the images, while still being a convex functional. The additional TV regularization of the motion field restricts the solution space and favors the piecewise constant fields typical for moving objects in a scene. Finally, the joint formulation is not only more elegant, but increases efficiency and accuracy.

In our experiments we show that this yields better results than both the pointwise optimization on alternate exposure images [23] and the state-of-the-art TV- L^1 approach on short exposure images [27]. We also provide results of real world recordings, as commercially available video cameras can be readily programmed to alternate between different exposure times and gain settings for every frame.

2 Related Work

The number of articles on computing the optical flow is tremendous, which indicates the significance of the problem as well as its severity [1, 5, 19]. Related to our work, scale-space approaches obtain reliable optical flow results in the presence of disparities larger than a few pixels [7, 22]. Alternatively, Lim *et al.* circumvent the problem by making use of high-speed camera recordings [18]. Alvarez *et al.* determine occluded regions by calculating forward and backward optical flow and checking for consistency [2]. Areas with large optical flow discrepancies are considered occluded and are simply excluded from further computations. Xiao *et al.* propose interpolating motion into occluded areas from nearby regions by bilateral filtering [25]. This ap-

proach is refined by Sand and Teller [22] in the context of particle video.

There has been some previous work on calculating motion fields from a single, motion-blurred image based on Fourier analysis [21] or auto-correlation [20] assuming spatially invariant motion. A recent approach [12] is able to calculate parametric and non-parametric motion fields formulating a constraint on the alpha channel of the blurred image. Motion estimation from a single motion-blurred image is one step of blind image deblurring approaches. Because deconvolution is in general ill-posed, these approaches are usually restricted to spatially invariant point spread functions (PSF) [14, 15, 16] or a locally invariant PSF [17]. To simplify the problem of blind image deblurring, many approaches use additional images to estimate the motion and to reconstruct the image: Yuan *et al.* use pairs of blurred and noisy images not only to estimate the spatially invariant blur kernel but also to reduce ringing artifacts during deconvolution [26]. The hybrid camera of Ben-Esra and Nayar takes a long-time exposure of the scene, while a detector with a much lower spatial and a higher temporal resolution takes a sequence of short-time exposures to detect camera motion [6]. From the camera motion, a global PSF can be reconstructed which is used to deblur the image. A recent extension of the hybrid camera [24] permits the kernel to be a local mixture of predefined basis kernels, which can be handled by modern deblurring methods. The motion-from-smear approach [10] focuses on motion detection from two motion-blurred images, using deconvolution techniques and thus relying on locally constant motion. In an extension [9], a short and a long exposed image are used to calculate the parameters of an affine motion model. The approach of Bar *et al.* considers two motion-blurred images to segment an

image into static background and a foreground that moves with constant velocity [4]. In a similar approach [13], at least two motion-blurred images are used to determine local motion, the corresponding segmentation and depth information of the scene restricting the motion to be a sideways translation parallel to the image plane.

Using short-long-short exposures, Fig. 1, our approach substantially improves on preliminary work [23]. Based on an image formation model that is able to handle occlusions as well as large displacements, the approach estimates dense motion fields directly without previous deblurring. In this paper, we investigate total variation regularized optimization instead of pointwise optimization, yielding considerably improved optical flow results. We can propagate motion information into textureless regions and achieve robustness against noise, a major limitation of the previous approach. By embedding the resulting energy formulation into the duality-based approach of [27], we show that exchanging an intermediate short exposure for a motion-blurred image improves dense motion field estimation.

3 Image Formation Model

In order to exploit the information provided by the long-time exposed image, we review the image formation model of [23] that relates the acquired images via a dense 2D motion field. As input, we assume two short exposure images $I_1, I_2 : \Omega \rightarrow \mathbb{R}$ which are taken before and after the exposure time of a third, long exposure input image $I_B : \Omega \rightarrow \mathbb{R}$. We look for a description of a motion blurred image $B : \Omega \rightarrow \mathbb{R}$ in terms of I_1, I_2 and the motion.

To simplify the model, we assume that the short exposures are free of motion blur, that the short and long exposed images are brightness-adjusted such that in case of no motion, all images are identical. We assume that scene surface appearance does not change considerably over time.

3.1 Without Occlusion

In the simplest case, one considers a moving scene without any occluded or disoccluded scene points, which implies that all scene points contributing to B are visible in I_1 as well as I_2 . Parametrizing by

time $t \in [0, 1]$ we obtain

$$B(x) = \int_0^1 I_1(\mathbf{p}_1(x, t)) dt = \int_0^1 I_2(\mathbf{p}_2(x, t)) dt, \quad (1)$$

where $\mathbf{p}_1(x, \diamond) : [0, 1] \rightarrow \Omega$ and $\mathbf{p}_2(x, \diamond) : [0, 1] \rightarrow \Omega$ are spatially varying, planar curves on the image plane, Fig. 2. In the case without occlusion, the entire curves are visible in both images, so that the values of both integrals are equal.

For ease of computation, we adopt a linear model for the motion curves:

$$\mathbf{p}_1(x, t) = x - t \mathbf{w}_1(x) \quad \text{and} \quad \mathbf{p}_2(x, t) = x + t \mathbf{w}_2(x), \quad (2)$$

where $\mathbf{w}_i : \Omega \rightarrow \mathbb{R}^2$. This turns out to be a sufficient approximation also for different types of motion considered in the experimental section. However, it should be noted that the model and algorithm are not inherently limited to the linear case.




Figure 2: Without occlusion, a contiguous path of scene surface points between y_1 and y_2 contributes to pixel x in the long-exposure image. The projection of the path onto the image plane forms a planar curve in the preceding (a) and in the succeeding (b) short exposure. For an occluded point the path is no longer contiguous and only partly visible in the projection (c).

3.2 With Occlusion

The long exposed image also permits incorporating occlusion effects into the image formation model. One assumes that a point changes its visibility at most once. If a scene surface becomes occluded, the integral is partitioned so that part of the brightness $B(x)$ observed in x is due to brightness along curve \mathbf{p}_1 , while the remaining part is due to brightness along \mathbf{p}_2 ,

$$B(x) = \int_0^s I_1(\mathbf{p}_1(x, t)) dt + \int_0^{1-s} I_2(\mathbf{p}_2(x, t)) dt. \quad (3)$$

Here, $s = s(x) \in [0, 1]$ denotes the moment during exposure where an object previously visible at x in I_1 becomes occluded by an object visible at x in I_2 , or vice versa.

Note that in the case of no occlusion, any choice of s yields the same brightness $B(x)$. One consequence is that s is not suitable to decide whether a point is occluded / disoccluded or remains visible. Instead, it is a mean to calculate truthful motion curves in occluded or disoccluded regions. The fact that s is indeterminate in the case of no occlusion has yet another consequence. One can then consider a fixed x and differentiate (3) with respect to s , arriving at the brightness constancy assumption of traditional optical flow computation. The alternate exposure approach is more general because it incorporates the additional information provided by the motion-blurred image I_B , and explicitly takes occlusion into account.

3.3 With Temporal Offset

We also allow for exposure gaps between the images I_1 and B as well as between B and I_2 . Gaps between exposures can occur, e.g., due to camera hardware constraints. To account for gaps, we include a temporal offset in (3), i.e., we change the integration limits by constants corresponding to the relative lengths of the gaps.

3.4 From Motion Curves to Displacement Fields

The motion curves considered describe the motion centered on the motion-blurred image. Since for many applications a displacement field is needed, we warp the motion curves \mathbf{p}_1 and \mathbf{p}_2 according to the estimated motion and occlusion parameters and thus obtain a displacement field for pixels in I_1 or I_2 , respectively.

4 Energy Formulation

The image formation model for a motion-blurred image B considered in the previous section provides us with a pointwise error measure for estimates of $\mathbf{w}_1, \mathbf{w}_2$ and s as follows. Given two short exposures I_1, I_2 and a long exposed image I_B , we can compare the blurred image I_B , i.e. the actual measurement to the result B predicted by the

model. We choose to consider the differences of the images in L^1 norm, as this norm is more robust against outliers in intensity values as the L^2 norm considered in the approach of [23].

$$\rho_1(\mathbf{w}_1, \mathbf{w}_2, s) = |I_B(x) - B(x)|. \quad (4)$$

For the sake of increased subpixel accuracy we also consider a differentiated version, i.e. the brightness constancy assumption

$$\rho_2(\mathbf{w}_1, \mathbf{w}_2, s) = |I_1(x - s\mathbf{w}_1) - I_2(x + (1 - s)\mathbf{w}_2)|. \quad (5)$$

Integrating the weighted sum of the pointwise errors over the image domain, we obtain a data term

$$E_{data}(\mathbf{w}_1, \mathbf{w}_2, s) = \int_{\Omega} \rho_1 + \gamma \rho_2 \, dx. \quad (6)$$

Instead of minimizing the pointwise error, we can increase stability and performance in textureless regions by considering global relationships of scene movements: Neighboring points belonging to the same object typically exhibit similar motion. This observation suggests including a regularization term in the energy functional. As demonstrated in previous work, using the total variation as a regularizer for flow fields produces very desirable results [27]. It favors piecewise constant motion fields, thus smoothing out undesired outliers and avoiding oversmoothing at motion boundaries at the same time. We also regularize the occlusion time, as neighboring pixels, if they are occluded at all, are occluded at related instants in time. The resulting energy which depends on the unknown motion curves $\mathbf{w}_1, \mathbf{w}_2$ and occlusion time s can be written as

$$E_{data} + \int_{\Omega} \alpha \left(\sum_{i=1}^2 |\nabla w_{1,i}| + |\nabla w_{2,i}| \right) + \beta |\nabla s| \, dx. \quad (7)$$

Here, $\alpha, \beta \geq 0$ are free parameters of our method, controlling the desired smoothness of the flow fields and of the moment of occlusion, respectively.

4.1 Minimization Method

Our minimization scheme is based on the primal-dual algorithm used for TV- L^1 Optical Flow [27], whose variants currently rank in the top of the Middlebury benchmark [3]. We briefly review the method here and show how we use this framework to minimize our more complex energy functional in

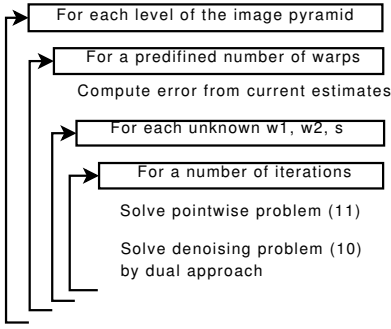


Figure 3: The workflow of our algorithm.

section 4.2 by replacing the general variable with w_1 , w_2 and s in turn and keeping the others fixed.

In the very general case that one wants to minimize a total variation energy of the form

$$E(\mathbf{u}) = \lambda \int_{\Omega} |\rho(\mathbf{u})| dx + \int_{\Omega} \sum_{i=1}^k |\nabla u_i| dx \quad (8)$$

for a k -dimensional function \mathbf{u} on Ω with a pointwise error term ρ , one can introduce an auxiliary vector field \mathbf{v} and instead consider the convex approximation

$$E_{\theta}(\mathbf{u}, \mathbf{v}) = \int_{\Omega} \lambda |\rho(\mathbf{v})| + \frac{1}{2\theta} \|\mathbf{u} - \mathbf{v}\|^2 + \sum_{i=1}^k |\nabla u_i| dx \quad (9)$$

If θ is small, \mathbf{v} will be close to \mathbf{u} near the minimum, and thus E will be close to E_{θ} . The key result of [27] is that the above energy can be minimized very efficiently using an alternating scheme, where one iterates between solving TV image denoising problems for each u_i , keeping \mathbf{v} fixed

$$\operatorname{argmin}_{u_i} \int_{\Omega} \frac{1}{2\theta} (u_i - v_i)^2 + |\nabla u_i| dx, \quad (10)$$

and a minimization problem for \mathbf{v} with fixed \mathbf{u}

$$\operatorname{argmin}_{\mathbf{v}} \lambda |\rho(\mathbf{v})| + \frac{1}{2\theta} \|\mathbf{u} - \mathbf{v}\|^2, \quad (11)$$

which can be solved point-wise with a thresholding scheme. Details and a proof of convergence can be found in [8, 27].

4.2 Application of the Framework

In our case, we use some slight modifications, adapted to our problem of minimizing the energy in terms of w_1 , w_2 and s . First, we employ the above scheme, i.e. iterating between (10) and (11), with $\mathbf{u} = w_1$, $\mathbf{u} = w_2$ or $\mathbf{u} = s$, respectively, to solve for each of our unknowns, given a fixed approximation of the two others. Then, if we use only the pointwise error $\rho = \rho_1$ in (11) we can directly apply the thresholding scheme detailed in [27]. Yet, if we want to incorporate ρ_2 , this scheme is no longer directly applicable and we therefore apply a descent scheme for (11), substituting the L^1 norm with its regularized variant $|\rho_i|_{\epsilon} = \sqrt{|\rho_i|^2 + 0.001}$.

4.3 Implementation

Since the techniques for the actual minimization are well known, we only point the reader towards the references where details on the numerical implementation of each step can be found [27].

In order to speed up convergence, we implement the algorithm on a scale pyramid of factor 0.5, initializing with $s = 0.5$ for occlusion timings, and zero motion curves on the coarsest level. On each level of the pyramid we perform several warping iterations where in each iteration we solve for s , w_1 and w_2 . For each variable an instance of (10) and (11) has to be solved, Fig. 3.

For (10), we employ the dual formulation detailed in [27], Prop. 1, using 5 iterations and a time step of $\tau = 0.1225$.

Suitable values for the parameter α , β , γ and θ were found experimentally.




Figure 4: The color code used to display motion fields in Fig. 5 and Fig. 6.



Figure 5: Scenes *Ben*, *windmill* and *corner*: Input images I_1 , I_B and I_2 (a)-(c), motion field calculated with our method (d).

5 Experiments

To evaluate our method we consider synthetic images as well as real-world recordings: We calculate motion fields for synthetic scenes with known ground-truth motion fields and compare the mean angular error (MAE) of the motion field with related approaches [23, 22, 27]. Note that we cannot evaluate our method on standard test data because the blurred image is not available. We also show results for real world recordings. The recordings were made with a commercially available Point-Grey Flea2 camera that is able to acquire short and long exposed images alternatingly.

For all experimental results we use a 5-level image pyramid, 10 warping iterations and 10 iterations to solve Eqs. (10) and (11). For normalized intensity values we found $\theta \in (0, 1]$, $\alpha, \beta \in (0, 0.1]$ and $\gamma \in [0, 0.5]$ to be suitable parameter values.

5.1 Synthetic Test Scenes

We consider synthetic test scenes containing different kinds of motion. The scene *Ben* contains translational motion of up to 14 pixels, the scene *windmill* depicts rotational motion approximately par-

allel to the image plane, while the scene *corner* exhibits rotation of objects around a vertical axis, Fig. 5. In all three cases, many images were rendered at short time intervals and averaged to obtain the motion-blurred images I_B . The first and the last rendered image represent the short-time exposures I_1 and I_2 .

To evaluate the advantage of the global optimization framework, we compare the results of our algorithm to the results of the pointwise algorithm of [23]. We also compare to state-of-the-art optical flow algorithms. For fair comparison, besides images I_1 and I_2 we provide the competing optical flow algorithms with the image $I_{1.5}$, depicting the scene half way between I_1 and I_2 . We calculate the motion fields between I_1 and $I_{1.5}$ as well as between $I_{1.5}$ and I_2 . The two results are then concatenated before comparing them to the ground truth displacement field. As optical flow works best for small displacements, the error of the concatenation is smaller than calculating the motion field between I_1 and I_2 directly.

For comparison, we consider two different optical flow methods. We chose the algorithm of Zach *et al.* [27], since it relies on the same math-

	<i>Ben</i>		<i>windmill</i>		<i>corner</i>	
	MAE	STD	MAE	STD	MAE	STD
Sand, Tellers [22]	8.42	20.91	6.78	17.43	6.40	17.71
Zach <i>et al.</i> [27]	5.81	20.08	4.87	17.35	5.05	19.59
Sellent <i>et al.</i> [23]	6.31	19.53	8.64	23.70	12.87	27.69
our method	4.27	16.35	4.56	15.70	4.57	17.32

Table 1: The motion field computed with our method has a smaller mean angular error (MAE) and a smaller standard deviation (STD) than motion fields computed with competitive optical flow or alternate exposure algorithms.




Figure 6: Real-world recordings *juggling* (top row) and *waving* (bottom row): I_1 , I_B , I_2 (a)-(c) and the motion field calculated with our method (d).

emational framework as our approach. However, our method uses a long exposed image instead of a higher frame rate of short exposed images. We also compare to the algorithm of Sand and Teller [22] on three images, as both our method and their approach consider occlusion effects while calculating motion fields. As can be seen in Table 1, our algorithm performs best in all three test scenes: the mean angular error of the estimated motion fields is smaller, as well as its standard deviation.

5.2 Real-World Recordings

We also test our method on real-world recordings. We use the built-in HDR mode of a PointGrey Flea2 camera to alter exposure time and gain between successive frames. By adjusting the gain, we ensure that corresponding pixels of static regions in the short and long exposed images are approximately of same intensity. With the HDR mode we are able to acquire I_1 , I_B and I_2 with a minimal time gap

between the images. The remaining gap is due to the fix 30 fps camera frame rate and the readout time of the sensor. The recorded images and the estimated motion fields are shown in Fig. 6. The *juggling* scene demonstrates vividly the advantages of using short and long exposures: the motion is very fast and sharp images require short exposure times of 6.02 ms. Yet the camera can only process an image every 33.33 ms. This leads to long gaps of 27.31 ms of unrecorded motion between sharp images. For our method, we record a long exposed image with an exposure time of 39.65 ms reducing the gap between I_B and the succeeding short exposed image to 0.48 ms and providing us with additional information. While the short exposures either show the ball or not, the motion-blurred image captures the path taken by the ball and enables correct motion field estimation, i.e. our method can handle the small ball leaving the picture.

For the *waving* scene we use exposure times of

20.71 ms and 124.27 ms, resulting in gaps of 12.45 ms and 0.48 ms, respectively. Note that the algorithm is capable of dealing with disoccluded texture as in the *waving* scene where the hand uncovers the face.

5.3 Limitations

Our method shares some of the limitations inherent to all optical flow methods which also the use of a motion-blurred image cannot remedy. Like in all purely image based methods, motion in poorly textured regions cannot be detected uniquely, as can be seen in the black background of the *waving* scene, Fig. 6. Also common to all optical flow methods, we assume that motion is the only source of change in brightness, disregarding highly reflecting and transparent surfaces from the calculations. In contrast to most optical flow methods, however, we include occlusion explicitly into our image formation model. We limit each pixel to change visibility only once during the exposure interval. This assumption is justified if the frame rate is sufficiently high. Several changes in visibility would partition the motion curves into more than two parts, thus rendering the problem numerically unstable.

6 Conclusion and Future Work

In this work, we proposed a variational approach to optical flow from a set of alternately exposed images. We refine the optimization of a general image formation model, that is able to handle occlusions, large displacements and objects moving out of the image. Based on the image formation model, we derive a TV- L^1 energy functional which we solve with an efficient dual method. In the experiments we show that making use of a long exposure image improves the accuracy of the motion field calculation: the mean angular error and its standard deviation is reduced, not only for linear motion but also for rotational types of motion. So far, we model changes in illumination only implicitly by the robust L^1 norm, but we hope to incorporate this into future motion-blurred image formation models. In future work, we also want to apply the optimization method to non-linear motion paths which promises to be favorable for frame interpolation.

Acknowledgements

The authors gratefully acknowledge funding by the German Science Foundation from project DFG MA2555/4-1.

References

- [1] J. Aggarwal and N. Nandhakumar. On the computation of motion from sequences of images - a review. In *Proc of the IEEE*, volume 76, 1988.
- [2] L. Alvarez, R. Deriche, T. Papadopoulos, and J. Sanchez. Symmetrical dense optical flow estimation with occlusions detection. *IJCV*, 75(3):371–385, 2007.
- [3] S. Baker, D. Scharstein, J. Lewis, S. Roth, M. Black, and R. Szeliski. A Database and Evaluation Methodology for Optical Flow. In *Proc ICCV*, 2007.
- [4] L. Bar, B. Berkels, M. Rumpf, and G. Sapiro. A variational framework for simultaneous motion estimation and restoration of motion-blurred video. In *Proc ICCV*, pages 1–8, 2007.
- [5] J. Barron, D. Fleet, and S. Beauchemin. Performance of optical flow techniques. *IJCV*, 12(1):43–77, 1994.
- [6] M. Ben-Ezra and S. K. Nayar. Motion deblurring using hybrid imaging. In *Proc CVPR*, pages 657–664, 2003.
- [7] T. Brox, A. Bruhn, N. Papenbergh, and J. Weickert. High accuracy optical flow estimation based on a theory for warping. In *Proc. ECCV*, pages 25–36, 2004.
- [8] A. Chambolle. An algorithm for total variation minimization and applications. *J Math Imag Vis*, 20(1):89–97, 2004.
- [9] W.-G. Chen, N. Nandhakumar, and W. N. Martin. Estimating image motion from smear: a sensor system and extensions. In *Proc ICIP*, pages 199–202, 1995.
- [10] W.-G. Chen, N. Nandhakumar, and W. N. Martin. Image motion estimation from motion smear—a new computational model. *IEEE T-PAMI*, 18(4), Apr. 1996.
- [11] W. Christmas. Filtering requirements for gradient-based optical flow measurement. *IEEE T-IP*, 9:1817–1820, Oct 2000.
- [12] S. Dai and Y. Wu. Motion from blur. In *Proc CVPR*, pages 1–8, 2008.
- [13] P. Favaro and S. Soatto. A variational approach to scene reconstruction and image segmentation from motion-blur cues. In *Proc CVPR*, 2004.
- [14] R. Fergus, B. Singh, A. Hertzmann, S. Roweis, and W. Freeman. Removing camera shake from a single photograph. In *SIGGRAPH*, pages 787–794, 2006.
- [15] J. Jia. Single image motion deblurring using transparency. In *Proc CVPR*, pages 1–8, 2007.
- [16] D. Kundur and D. Hatzinakos. Blind image deconvolution. *IEEE Signal Process Mag*, 13(3):43–64, 1996.

- [17] A. Levin. Blind motion deblurring using image statistics. *Adv Neural Inform Process Syst*, 19:841–848, 2007.
- [18] S. Lim, J. Apostolopoulos, and A. Gamal. Optical flow estimation using temporally oversampled video. *IEEE T-IP*, 14(8):1074–1087, 2005.
- [19] A. Mitiche and P. Bouthemy. Computation and analysis of image motion: A synopsis of current problems and methods. *IJCV*, 19(1):29–55, 1996.
- [20] T. Pao and M. Kuo. Estimation of the point spread function of a motion-blurred object from autocorrelation. In *Proc of SPIE*, volume 2501, 2003.
- [21] I. M. Rekleitis. Optical flow recognition from the power spectrum of a single blurred image. In *Proc ICIP*, pages 791–794, 1996.
- [22] P. Sand and S. Teller. Particle video: Long-range motion estimation using point trajectories. pages 2195–2202, 2008.
- [23] A. Sellent, M. Eisemann, and M. Magnor. Motion Field and Occlusion Time Estimation via Alternate Exposure Flow. In *IEEE International Conference on Computational Photography*. IEEE, April 2009.
- [24] Y. Tai, H. Du, M. Brown, and S. Lin. Image/video deblurring using a hybrid camera. In *Proc CVPR*, pages 1–8, 2008.
- [25] J. Xiao, H. Cheng, H. Sawhney, C. Rao, and M. Isnardi. Bilateral filtering-based optical flow estimation with occlusion detection. In *Proc ECCV*, pages 211–224, 2006.
- [26] L. Yuan, J. Sun, L. Quan, and H.-Y. Shum. Image deblurring with blurred/noisy image pairs. In *SIGGRAPH*, number 3, 2007.
- [27] C. Zach, T. Pock, and H. Bischof. A duality based approach for realtime TV- L^1 optical flow. In *Pattern Recognition*, volume 4713, pages 214–223, 2007.