

Coded Aperture Flow

Anita Sellent* and Paolo Favaro

Institut für Informatik und angewandte Mathematik,
Universität Bern, Switzerland
<http://www.cvg.unibe.ch/>

Abstract. Real cameras have a limited depth of field. The resulting defocus blur is a valuable cue for estimating the depth structure of a scene. Using coded apertures, depth can be estimated from a single frame. For optical flow estimation between frames, however, the depth dependent degradation can introduce errors. These errors are most prominent when objects move relative to the focal plane of the camera. We incorporate coded aperture defocus blur into optical flow estimation and allow for piecewise smooth 3D motion of objects. With coded aperture flow, we can establish dense correspondences between pixels in succeeding coded aperture frames. We compare several approaches to compute accurate correspondences for coded aperture images showing objects with arbitrary 3D motion.

Keywords: Coded Aperture, Optical Flow

1 Introduction

Optical flow algorithms estimate the apparent motion between succeeding frames of a video sequence [6] by comparing the brightness values of pixels. Optical flow is an approximation of the projection of 3D motion to the image plane. Traditionally, optical flow algorithms consider pinpoint sharp images, without any degradations other than moderate levels of noise [3]. Real recording conditions, however, rarely allow to capture pinpoint sharp images. When the amount of light in a scene is limited, real cameras require a finite aperture to capture images with a usable signal to noise ratio. Finite aperture sizes introduce defocus blur into images of non fronto-planar scenes. We found that this depth dependent image degradation can lead to erroneous optical flow estimates. However, the size of the blur provides depth information. In fact, defocus blur is a frequently exploited depth cue [18].

Conventional depth from defocus approaches acquire several images of a static scene to estimate a depth map and reconstruct sharp textures. By introducing a coded mask into the aperture of a conventional camera, depth estimates as well as texture restoration can be obtained from a single input image [9, 21]. This single image method is highly suited to provide monocular depth cues in

* Funded by the Deutsche Forschungsgemeinschaft, project Se-2134/1

dynamic scenes, where 3D location and shape of every object in a scene changes independently from frame to frame.

For the estimation of pixel trajectories over time, coded aperture frames are a challenging input. The appearance of objects changes dramatically whenever they move relative to the focal plane. Conventional optical flow algorithms do not take this change into account. In this work we consider several approaches to model the effect of defocus blur in optical flow estimation. We evaluate these formulations in the particular setup of high frequency aperture masks that are optimized for the estimation of depth from a single input frame.

2 Related Work

The estimation of optical flow from image sequences is a challenging problem. For a summary and evaluation of modern approaches we refer the reader to the work of Sun et al. [19] and Baker et al. [3]. In our work we build upon the TV- L^1 optical flow approach of Zach et al. [23] and its anisotropic extension by Werlberger et al. [22]. These approaches estimate dense optical flow with a robust L^1 norm for comparing brightness values in two frames and (anisotropic) total variation regularization. The algorithms use a dual optimization scheme that, as a GPU implementation, allows for real time dense flow estimation with state-of-the-art accuracy. While most optical flow algorithms ignore depth altogether, some approaches assign each pixel to a layer and can thus achieve improved regularization and high accuracy, see e.g. Ref. [20]. Still, the layers do not incorporate a model for defocus blur that may change from frame to frame. In contrast, the filter flow by Seitz and Baker [14] models *relative blur* between two images, and allows to compute accurate correspondences also in the presence of defocus blur. For the high frequency apertures that enable single frame depth estimation, relative blurring is not applicable. Thus for coded aperture flow deblurring is necessary for the comparison of points that move in depth. In section Sec. 3 we adapt the filter flow to coded apertures to evaluate its performance in our application.

Other approaches to consider defocus effects in dynamic scenes have been introduced by Kubota et al. [8] and Shroff et al. [17]. Both build on the assumption that objects do not move in depth. Shroff et al. acquire focal stacks of moving scenes. They initialize optical flow estimation on images with the same focus settings. Then they refine this flow by considering images with different focal settings, re-blurring deblurred images according to the current, constant depth estimate. Kubota et al. avoid deblurring by applying blur to both images. They evaluate all possible combinations of blur for the best correspondences, before a common depth map is estimated with a depth-from-defocus approach. In contrast to this approach, coded apertures allow to estimate depth from a single frame. In our setup we can therefore simplify correspondence estimation by applying the estimated depth map directly.

Finite apertures improve signal to noise ratio by admitting more light than the ideal pinhole camera. When instead the exposure time is extended, the im-

ages are affected by motion blur. The modeling of motion blur can improve the performance of tracking [7] or dense optical flow estimation [13]. However, extended exposure time does not introduce additional information on the depth structure of a dynamic scene such as provided by the coded apertures. To obtain depth information Xu and Jia consider stereo images and then remove depth dependent motion blur. However, even for the stereo camera setup, scene motion is restricted to mainly translational camera motion. Independent object motion is not allowed.

In our work we profit from defocus blur as a depth cue. There are some recent advances in single image depth estimation using a conventional aperture, e.g. Ref. [10]. However, using a conventional aperture sacrifices high frequency image content in the low-pass property of the full-aperture blur. Coded aperture images preserve high frequency content more faithfully. This preservation property can be used for improved deblurring results. By evaluating the quality of the images deblurred with different depth hypotheses, a depth map can be estimated, [9,21]. In contrast, filter-based coded aperture depth estimation [5, 11] evaluates the blurred images themselves for depth estimation. More accurate results can be obtained with less computational burden. To obtain smoother and temporally consistent depth maps, Martinello and Favaro [12] use succeeding frames in a coded aperture image sequence for regularization of the depth map. However, they do not compute explicit correspondences between frames and exploit only objects that move parallel to the image plane. We focus on scenes where objects can move arbitrarily and estimate their motion from frame to frame.

3 Brightness Constancy Assumptions

The basic assumption of optical flow estimation is that the brightness of a pixel does not change through the motion [6]. Given two focused images $I_1, I_2 : \Omega \rightarrow [0, 1]$ of a moving scene, this brightness constancy assumption can be written as $I_1(x) = I_2(x + u)$ for $x \in \Omega \subset \mathbb{R}^2$ and a displacement vector $u \in \mathbb{R}^2$. To solve this equation, usually the Taylor linearization $I_2(x + u) \approx I_2(x) + u^\top \nabla I_2(x)$ is used. For focused images, the estimation of the displacement u is then based on the data-term

$$D_F(x, u) = I_2(x) + u^\top \nabla I_2(x) - I_1(x) \quad . \quad (1)$$

When a planar scene is out of focus, we measure the defocused image $B_1 = k_1 * I_1$ where the defocus blur is expressed as the convolution with a depth dependent point spread function (PSF) k_1 . In this case, the brightness of a pixel x depends on the PSF and the brightness of the neighboring pixels. The PSF is depth dependent, so in sum the brightness of a pixel depends on the neighborhood and its depth. When a surface point moves towards the camera, its brightness in every frame is different.

In coded aperture photography, aperture masks are designed to effect a texture highly differently for different depths [15]. For optical flow estimation based on the brightness constancy assumption, highly depth dependent brightness is a hinderance. On the other hand, the coded aperture masks allow for single frame

depth estimation. We can profit from the estimated depth to improve optical flow estimation.

In our evaluation, we concentrate on the estimation of optical flow on coded aperture images. For depth estimation we use the state-of-the-art algorithm of Martinello and Favaro [11]. In the following we assume that a spatially variant depth map $d_i : \Omega \rightarrow \mathbb{R}$ is given for all measured frames B_i .

We compare several approaches to obtain optical flow estimates for coded aperture images. The first approach is based on the idea that high frequency aperture masks conserve high image frequencies better than conventional apertures. They therefore provide better deblurring results [9]. The deblurred images \hat{I}_1, \hat{I}_2 are all-in-focus representation of the scene, for which the linearized brightness constancy, Eq. (1), can be directly applied

$$D_D(x, u) = \hat{I}_2(x) + u^\top \nabla \hat{I}_2(x) - \hat{I}_1(x) \quad (2)$$

Thereby we compute the images \hat{I}_i by the conjugate gradient based, spatially variant deblurring that has been applied successfully for coded aperture images with multiple objects [16]. We use the optimized smoothness weight of 0.01 and 50 iterations. The advantage of operating optical flow estimation on the deblurred image is that any state-of-the-art optical flow implementation can be used out-of-the-box. The disadvantage is that we need to perform the ill-posed procedure of deblurring twice. Apart from disturbance by deblurring artifacts, we also have to deal with the computational burden of the deblurring.

Under the assumption of fronto-parallel scene patches we also consider a second approach by adapting the idea of Refs. [14,17] to coded aperture images. We compare the measured image B_2 to the re-blurred image $k_{d_2} * \hat{I}_1$ where k_{d_2} is the PSF corresponding to the estimated depth in $B_2(x+u)$. Under the assumption of local planarity, we can apply Taylor linearization and obtain the brightness term

$$D_S(x, u) = B_2(x) + u^\top \nabla B_2(x) - (k_{d_2(x)} * \hat{I}_1)(x) \quad (3)$$

In comparison to the term D_D in Eq. (2) we now only have to estimate the deblurred image \hat{I}_1 . Still, this procedure can introduce deblurring artifacts that might not be compensated by the convolution with k_{d_2} .

Under the same assumptions as above, we also consider a third approach. Based upon the idea of mutual blurring of Refs. [7,8,13] we compare the images $k_{d_1} * B_2$ and $k_{d_2} * B_1$. Linearization leads to the brightness term

$$D_M(x, u) = (k_{d_1(x)} * B_2)(x) + u^\top \nabla (k_{d_1(x)} * B_2)(x) - (k_{d_2(x)} * B_1)(x) \quad (4)$$

This approach has the advantage that deblurring is not required. However, relying on mutual blurring of defocused images potentially sacrifices those image frequencies that are required for accurate optical flow estimation.

From the three above formulations of depth dependent brightness constancy we want to evaluate which provides us with the most accurate flow estimates.

Generally, scenes consist of multiple objects and therefore incorporate depth discontinuities. The brightness of pixels at depth discontinuities is determined by objects in the foreground as well as the background [2]. In this case convolution with a single, depth dependent PSF is not an accurate description. Instead of introducing more elaborate defocus models, we decided to disable the brightness constancy assumption for pixels close to discontinuities. We introduce a weight function $\phi_{d_1}(x) = \exp(\frac{\int_{N_x} \|\nabla d_1(x)\|^2 dx}{\sigma_d})$ where N_x is a small neighborhood around pixel x and σ_d a constant. This weight function disables the brightness constancy assumption at known depth discontinuities, i.e., when ∇d_1 is large. We set N_x to $\frac{1}{4}$ the maximally considered blur size and fix $\sigma_d = 3$.

A further cue to depth discontinuities are occlusion boundaries of moving objects. As proposed by Alvarez et al. [1] we therefore compare forward motion estimate w and backward motion estimate $v : \Omega \rightarrow \mathbb{R}^2$. When the difference is large a point is most probable occluded. We let $\phi_s(x) = \exp(\frac{\|w(x) - v(x+w)\|^2}{\sigma_s})$ with the parameter σ_s set to 2% of the smallest image dimension. Our final confidence in the brightness constancy is $\phi_w = \phi_s \phi_d$.

4 Estimating Correspondence Fields

All brightness constancy assumptions introduced in the previous section provide only one equation for the two unknown components of the displacement vector. To solve for a dense optical flow field $w : \Omega \rightarrow \mathbb{R}^2, x \mapsto u = \begin{pmatrix} w_1 \\ w_2 \end{pmatrix}$, we additionally assume piecewise smoothness of the flow. A typical difficulty is the determination of the pieces to impose smoothness on. In conventional optical flow estimation, only the images are available to determine regions. In coded aperture flow we also have the depth map available. We expect the flow to be discontinuous at the same locations where the depth of the scene changes rapidly.

To study the effect of coded aperture blur in optical flow estimation on a comparable basis, we incorporate the modified brightness constancies and the depth dependent regularization in the state-of-the-art optical flow of Werlberger et al. [22]. For completeness, we here give a short summary of the approach highlighting our modifications. For more details on the original optical flow algorithm, we refer the reader to Ref. [22].

The first modification in our implementation is the data term. Instead of the conventional brightness constancy, Eq. 1, we consider alternative expressions, Eqs. (2) - (4). Additionally, we include an occlusion weight ϕ_w to circumvent false brightness comparisons at object boundaries. The second modification is to consider depth gradients for regularization. For the depth map normal $n_i = \frac{\nabla d_i}{\|\nabla d_i\|}$ and its perpendicular vector n^\perp we consider the diffusion tensor $T = \exp(a\|\nabla d_i\|)nn^\top + n^\perp(n^\perp)^\top$. Thus, our variational formulation of the problem takes the form

$$\min_{w: \Omega \rightarrow \mathbb{R}^2} \sum_{x \in \Omega} \lambda \phi_w |D_q| + \sum_{i=1}^2 \psi_\epsilon ((\nabla w_i)^\top T \nabla w_i) \quad (5)$$




Fig. 1: For the evaluation of coded aperture flow we render coded aperture frames for scenes of which the 3D motion, i.e. depth maps for each frame and the 2D projection of the motion, is known. From top to bottom: Input frames B_1 and B_2 and ground truth 2D motion, color coded with the map in Fig. 4c

with D_q either of our brightness constancy formulations D_D , D_S or D_M , $\lambda > 0$ a constant and ψ_ϵ the Huber-norm from Ref. [22]. Given the linearizations in Sect. 3 we can apply the solution scheme of Ref. [22]. For comparability we picked suitable parameters of the algorithm for conventional optical flow estimation and kept them fixed for all experiments. In detail, for normalized images we set $a = 0.20$, $\lambda = 50$, and, from Ref. [22] $\epsilon = 0.1$, $\theta = 1$ in the solution scheme.

The implementation of Werlberger et al. works on a Gaussian image pyramid to increase speed and obtain robust estimates for large displacements. To compute D_S and D_M for downsampled images, we require corresponding PSFs and depth maps. We downscale the PSF from the camera calibration, Sect. 5 to obtain PSFs for each level of the image pyramid. To obtain a down-sampled depth map, we consider all depth levels that contribute to a pixel on a coarser level and pick the depth level that is closest to the camera. This heuristic is motivated by the fact that for constant motion the projection of foreground motion spans larger 2D displacements. In our experiments we use 6 levels of an image pyramid with a down-sampling factor of 2.




Fig. 2: Estimating conventional optical flow on defocused images with objects moving in depth (scene *Wall*, Fig. 1) leads to erroneous flow estimation (a), (b). Deblurring the defocused input image with the estimated depth map provides visually pleasing images (c). Still, optical flow estimation between two deblurred images is noisy (d). Better results can be obtained when only one image is deblurred (e) or images are mutually blurred (f) (color coding with Fig. 4c).

5 Experiments

We evaluate the different approaches to calculate optical flow on coded aperture frames in several experiments. First we perform evaluation on synthetic images with known ground truth. Then we show results on real images.

5.1 Synthetic Experiments

We render several synthetic scenes with blur size between 4 and 11 pixels. The rendered frames for the 5×5 optimized coded aperture from Ref. [15] are shown in Fig. 1. The scenes contain different challenges ranging from a simple plane moving away from the focal plane, Fig. 1a, to a complex object moving in space, Fig. 1d. Note that for all experiments we keep all parameter of the algorithm fixed. All flow fields in this work are visualized with the color scale in Fig. 4c using black for points with $\phi_w(x) < 0.5$ that are rejected as occluded.

Accuracy Evaluation In our first experiment, we evaluate the different approaches to coded aperture flow for their accuracy. First we observe that optical flow estimation on defocused images with a conventional algorithm leads to noisy results, Figs. 2a and 2b. We also find that on synthetic images with known PSF and estimated depth map the results of the deblurring is visually very pleasing,

Table 1: We compare the average endpoint error of different formulations of brightness constancy. Computing optical flow (OF) on images with a conventional, full aperture in most cases results in a smaller error than optical flow on coded aperture images. Better results can be obtained when the estimated depth map is incorporated in the brightness constancy assumption by using D_D , D_S or D_M although the estimated depth has a certain mean squared error.

	OF, full	OF, coded	D_D	D_S	D_M	Depth
Wall	0.46 px	0.67 px	0.22 px	0.09 px	0.08 px	0.17 px ²
TriPlane	0.28 px	0.30 px	0.23 px	0.21 px	0.15 px	0.55 px ²
Slanted	0.68 px	0.85 px	0.49 px	0.10 px	0.06 px	0.20 px ²
Chair	0.58 px	0.61 px	0.36 px	0.38 px	0.28 px	0.65 px ²

Fig. 2c. Still, optical flow estimation between two deblurred images is noisy, Fig. 2d. Better results can be obtained by using re-blurred images or mutually blurred images, Figs. 2e, 2f. By evaluating the endpoint error of the unoccluded optical flow, Tab. 1, we observe that any formulation of depth dependent brightness improves the agnostic approach. The improvement is clearly visible, even though the estimated depth maps have a remaining depth estimation error. Over all our synthetic data-sets we observe best performance by the data-term based on mutual blurring. Although the deblurred images are visually pleasing, deblurring artifacts seem to deteriorate the accuracy of other approaches to coded aperture flow.

Similar results can be obtained for a variety of coded aperture masks proposed in literature, see supplementary material.

In our second experiment we evaluate the robustness of the coded aperture flow towards errors in the depth estimation. For our synthetic scenes, ground truth depth maps are known. For additional comparison we also use the point-wise depth estimates returned by the algorithm [11]. We compute flow fields with these depth maps as input and observe that deblurring both input images still gives the worst coded aperture flows, Tab. 2.

In the next experiment we evaluate the influence of the occlusion term. The effect is most prominent in the *Chair* sequence. E.g. for data-term D_M the average endpoint error for setting $\psi_w = 1$ is 1.43 px. By setting $\psi_w = \psi_d$, i.e. considering only the depth dependent cue, we can reduce the error to 0.97 px. Setting $\psi_w = \psi_s$ the error is reduced to 0.56 px. By combination of the terms with $\psi_w = \psi_s \psi_d$ a further reduction of the error to 0.28 px can be obtained (see supplement for the other data-sets).

Runtime Evaluation We implemented the coded aperture flow estimation using MATLAB. We use the same basic framework for each of the data-terms. Deblurring two images and estimating optical flow with data-term D_D takes 81 seconds on a 3.2GHz Mac Pro. Deblurring one image and employing data-term D_S takes 72 seconds. The deblurring free data-term D_M allows for optical flow estimation in 61 seconds.

Table 2: We evaluate the robustness of the different approaches to coded aperture flow towards the estimated depth map. Due to deblurring artifacts, D_D has the highest endpoint error even when ground truth depth is known (a). The point-wise estimated depth map is slightly less accurate than its smoothed version, but allows for comparable flow estimation, see (b) and Tab. 1

	(a)			(b)				
GT depth	D_D	D_S	D_M	Pointwise	D_D	D_S	D_M	Depth
Wall	0.14 px	0.06 px	0.07 px	Wall	0.20 px	0.09 px	0.08 px	0.18 px ²
TriPlane	0.10 px	0.06 px	0.04 px	TriPlane	0.25 px	0.21 px	0.15 px	0.58 px ²
Slanted	0.37 px	0.08 px	0.05 px	Slanted	0.47 px	0.10 px	0.06 px	0.21 px ²
Chair	0.12 px	0.12 px	0.11 px	Chair	0.34 px	0.34 px	0.26 px	0.68 px ²

5.2 Real Images

We acquire real image sequences by introducing the binary 5×5 mask from Ref. [15] into a Canon EF f/1.8 II lens [4]. We attach the lens to a Canon EOS 5D, Mark II camera that we set to continuous shooting mode. The camera is calibrated by acquiring a single point spread function (PSF) from a calibration point light source. The blur kernels for all other scales are generated synthetically from the measured image by downscaling to adjust to different depth levels.

Fig. 4 shows the scene *train*, and the optical flow we obtain with conventional algorithms and with coded aperture flow. Note how only the data-terms in Eqs. (3) and (4) can estimate the motion of the whole train correctly, even for the weakly textured locomotive. In the scene *walking* a person approaches the camera. Here all coded aperture approaches provide a good flow estimate in spite of noisy depth maps. However, deblurring both images introduces more noise on the background stones to the right than the other two approaches.

6 Conclusion

We consider dense optical flow between images that are acquired with a coded aperture. Unlike the ideal sharp image usually assumed for optical flow estimation, coded aperture defocus allows for single frames depth estimation. We show that conventional optical flow estimation is unsuitable to estimate accurate motion for objects moving relative to the focus plane. Instead, we evaluate three different formulations that take defocus maps into consideration for flow estimation. We find that the most accurate results can be obtained by comparing a measured image to a reblurred deconvolved image or by comparing mutually blurred images. As the latter approach is faster, we plan to use this approach in our future work on coded aperture video. Generally, the high accuracy that can be obtained with all evaluated methods also shows that coded aperture defocus blur preserves a sufficient amount of high frequency texture for dense optical flow estimation.

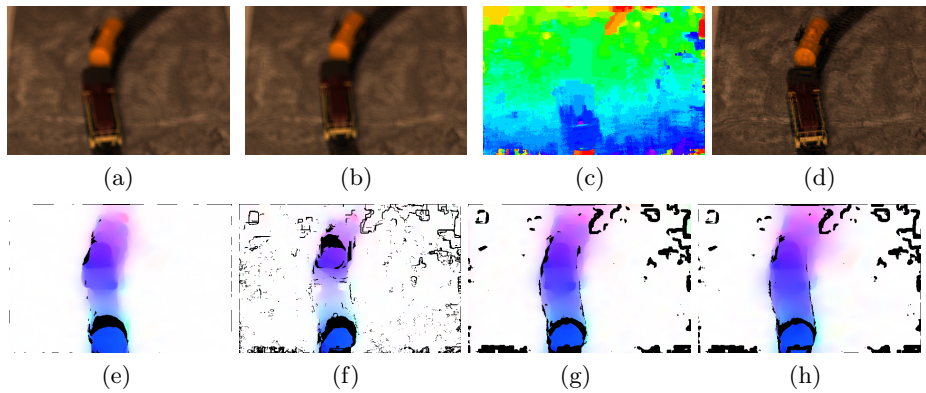


Fig. 3: A toy train backs away from the focal plane, (a) and (b). Coded apertures allow to estimate depth independently for each frame (c) and eases to deblur the images (d). Ignoring coded defocus effects in optical flow estimation (e) leads to inaccurate flow. Deblurring the images before conventional optical flow estimation (f) is susceptible to deblurring effects. Better results can be obtained by a combination of deblurring and re-blurring (g) or the application of mutual blur (h).




Fig. 4: A person approaches the camera, (a) Although the depth map is noisy (color coded with (d)) coded aperture flow estimation provides reasonable flow estimates (e) by deblurring both input images, (f) by re-blurring a deblurred image and (g) by applying mutual blur (color coded with (c)).

References

1. Alvarez, L., Deriche, R., Papadopoulos, T., Sánchez, J.: Symmetrical dense optical flow estimation with occlusions detection. In: *Computer Vision – ECCV 2002*, pp. 721–735. Springer (2002)
2. Asada, N., Fujiwara, H., Matsuyama, T.: Analysis of photometric properties of occluding edges by the reversed projection blurring model. *T-PAMI* 20(2), 155–167 (1998)
3. Baker, S., Scharstein, D., Lewis, J., Roth, S., Black, M., Szeliski, R.: A database and evaluation methodology for optical flow. *IJCV* 92(1), 1–31 (2011)
4. Bando, Y.: How to disassemble the canon EF 50mm f/1.8 II lens (2013), <http://web.media.mit.edu/~bandy/rgb/disassembly.pdf>
5. Dowski Jr, E., Cathey, W.: Single-lens single-image incoherent passive-ranging systems. *Applied Optics* 33(29), 6762–6773 (1994)
6. Horn, B.K., Schunck, B.G.: Determining optical flow. *Artificial intelligence* 17(1), 185–203 (1981)
7. Jin, H., Favaro, P., Cipolla, R.: Visual tracking in the presence of motion blur. In: *Proc. CVPR*. vol. 2, pp. 18–25. IEEE (2005)
8. Kubota, A., Kodama, K., Aizawa, K.: Registration and blur estimation methods for multiple differently focused images. In: *Proc. ICIP*. vol. 2, pp. 447–451 (1999)
9. Levin, A., Fergus, R., Durand, F., Freeman, W.: Image and depth from a conventional camera with a coded aperture. *TOG* 26(3), 70 (2007)
10. Lin, J., Ji, X., Xu, W., Dai, Q.: Absolute depth estimation from a single defocused image. *T-IP* 22(11), 4545–4550 (2013)
11. Martinello, M., Favaro, P.: Single image blind deconvolution with higher-order texture statistics. *Video Processing and Computational Video* pp. 124–151 (2011)
12. Martinello, M., Favaro, P.: Depth estimation from a video sequence with moving and deformable objects. *Proc. Image Processing Conference* (2012)
13. Portz, T., Zhang, L., Jiang, H.: Optical flow in the presence of spatially-varying motion blur. In: *Proc. CVPR*. pp. 1752–1759. IEEE (2012)
14. Seitz, S., Baker, S.: Filter flow. In: *Proc. ICCV*. pp. 143–150. IEEE (2009)
15. Sellent, A., Favaro, P.: Optimized aperture shapes for depth estimation. *Pattern Recognition Letters* 40, 96–103 (2014)
16. Sellent, A., Favaro, P.: Which side of the focal plane are you on? In: *Proc. ICCP*. pp. 1–8. IEEE (2014)
17. Shroff, N., Veeraraghavan, A., Taguchi, Y., Tuzel, O., Agrawal, A., Chellappa, R.: Variable focus video: Reconstructing depth and video for dynamic scenes. In: *Proc. ICCP*. pp. 1–9. IEEE (2012)
18. Subbarao, M., Surya, G.: Depth from defocus: a spatial domain approach. *IJCV* 13(3), 271–294 (1994)
19. Sun, D., Roth, S., Black, M.: Secrets of optical flow estimation and their principles. In: *Proc. CVPR*. pp. 2432–2439. IEEE (2010)
20. Sun, D., Wulff, J., Sudderth, E., Pfister, H., Black, M.: A fully-connected layered model of foreground and background flow. In: *Proc. CVPR*. pp. 2451–2458 (2013)
21. Veeraraghavan, A., Raskar, R., Agrawal, A., Mohan, A., Tumblin, J.: Dappled photography: Mask enhanced cameras for heterodyned light fields and coded aperture refocusing. *TOG* 26(3), 69 (2007)
22. Werlberger, M., Trobin, W., Pock, T., Wedel, A., Cremers, D., Bischof, H.: Anisotropic Huber-L1 optical flow. In: *Proc. BMVC*. pp. 1–11 (2009)
23. Zach, C., Pock, T., Bischof, H.: A duality based approach for realtime TV-L1 optical flow. In: *Pattern Recognition*, pp. 214–223. Springer (2007)