Postprocessing of Optical Flows Via Surface Measures and Motion Inpainting

Claudia Kondermann, Daniel Kondermann, and Christoph Garbe*

HCI at Interdisciplinary Center for Scientific Computing,
University of Heidelberg, Germany
claudia.kondermann@iwr.uni-heidelberg.de

Abstract. Dense optical flow fields are required for many applications. They can be obtained by means of various global methods which employ regularization techniques for propagating estimates to regions with insufficient information. However, incorrect flow estimates are propagated as well. We, therefore, propose surface measures for the detection of locations where the full flow can be estimated reliably, that is in the absence of occlusions, intensity changes, severe noise, transparent structures, aperture problems and homogeneous regions. In this way we obtain sparse, but reliable motion fields with lower angular errors. By subsequent application of a basic motion inpainting technique to such sparsified flow fields we obtain dense fields with smaller angular errors than obtained by the original combined local global (CLG) method and the structure tensor method in all test sequences. Experiments show that this postprocessing method makes error improvements of up to 38% feasible.

1 Introduction

Optical flow calculation is a crucial step for a wide variety of applications ranging from scientific data analysis and medical imaging to autonomous vehicle control and video compression. Despite high quality results on common test sequences, there are several challenging situations in image sequences, where many or even all known methods fail, e.g. in the case of difficult occlusions, transparent structures, severe noise, aperture problems, homogeneous regions or incoherent motion. Previously, some of these situations could be identified by means of confidence measures. In this paper we demonstrate that optical flow fields obtained by both local and global methods can be improved significantly by identification of situations, where a reliable estimation is possible, and subsequent motion inpainting in order to obtain a dense optical flow field. Our approach consists of two steps. First, we propose new confidence measures called "surface measures", which indicate the feasibility of a correct flow estimation. In a second step we minimize a basic motion inpainting functional in order to estimate the

* The authors thank the German Research Foundation (DFG) for funding this work within the priority program "Mathematical methods for time series analysis and digital image processing" (SPP1114).

G. Rigoll (Ed.): DAGM 2008, LNCS 5096, pp. 355–364, 2008.

flow within the unreliable regions only based on the flow information classified as reliable by the surface measures. In this way we obtain significantly improved optical flow fields. For the local structure tensor method we even obtain results superior to the global CLG method. Let a given image sequence $\mathcal{I}$ be defined on a space-time interval $\Omega \times [0,T]$, $\mathcal{I} : \Omega \times [0,T] \rightarrow \mathbb{R}$, $\Omega \subseteq \mathbb{R}^2$. Then the notion "optical flow" refers to the displacement field u of corresponding pixels in subsequent frames of an image sequence, $u : \Omega \times [0,T] \rightarrow \mathbb{R}^2$.

2 Related Work

Previously measures for the detection of difficult situations in the image sequence have been understood as a subclass of confidence measures. They mainly focus on the detection of features in the image, which make a correct estimation difficult. Such measures examine for example the magnitude of the image gradient or the eigenvalues of the structure tensor [1]. In contrast, other confidence measures are based on the flow computation method, e.g. [2] for variational methods or [3] for local methods. Limited comparisons of confidence measures have been carried out by Barron and Fleet [4] and Bainbridge and Lane [5]. Yet, the measures proposed so far are not able to detect all relevant difficult situations in an image sequence. Hence, to detect situations, where a reliable optical flow computation is feasible, we propose to examine the intrinsic dimensionality of image invariance functions similar to correlation surfaces. Correlation surfaces have for example been applied by Rosenberg and Werman [6] in order to detect locations where motion cannot be represented by a Gaussian random variable and by Irani and Anandan in [7] to align images obtained from different sensors. The inpainting of motion fields has been proposed before by Matsushita et al. [8] in order to accomplish video stabilization. Their approach differs from ours in two points: 1) The flow field is only extrapolated at the edges of the image sequence to continue the information to regions occluded due to perturbations of the camera. In contrast, we interpolate corrupted flow field regions within the sequence, which are identified by surface measures. 2) Instead of a fast marching method we use a variational approach to fill in the missing information.

3 Intrinsic Dimensions

According to [9] the notion 'intrinsic dimension' is defined as follows: 'a data set in n dimensions is said to have an intrinsic dimensionality equal to d if the data lies entirely within a d-dimensional subspace'. It has first been applied to image processing by Zetzsche and Barth in [10] in order to distinguish between edge-like and corner-like structures in an image. Such information can be used to identify reliable locations, e.g. corners, in an image sequence for optical flow computation, tracking and registration. An equivalent definition of 'intrinsic dimension' in image patches is based on its spectrum assigning the identifier

– i0d if the spectrum consists of a single point (homogeneous image patch)
– i1d if the spectrum is a line through the origin (edge in the image patch)
– i2d otherwise (e.g. edges and highly textured regions)

In [11,12] Barth has introduced the intrinsic dimensionality of three-dimensional image sequences and applied it to motion estimation, especially for multiple and transparent motions. Krüger and Felsberg [13] proposed a continuous formulation of the intrinsic dimension and introduced a triangular topological structure. Provided the assumption of constant brightness over time holds, motion of a single point corresponds to a line of constant brightness in the image sequence volume. Thus, the intrinsic dimension of locations in image sequences, where motion takes place is lower or equal to two. In the case of intrinsic dimension three the brightness constancy assumption is violated which can be due to e.g. noise or occlusions. Only in the i2d case a reliable estimation of the motion vector is possible, since then the trajectory of the current pixel is the only subspace containing the same intensity. Otherwise, severe noise (i3d), aperture problems (i1d) or homogeneous regions (i0d) prevent accurate estimates. As indicated in [11] occlusions or transparent structures increase the intrinsic dimension by one thus leading to the problem of occluded i1d locations misclassified as reliable i2d locations. Hence, the detection of i2d situations which do not originate from the occlusion of i1d situations would be beneficial for optical flow computation methods, as only in these cases reliable motion estimation is possible. This new situation will be denoted by 'i2d-o situation'.

4 Surface Measures

In order to estimate the accuracy of a given optical flow field we investigate the intrinsic dimension of invariance functions $f : \Omega \times [0,T] \times \mathbb{R}^2 \to \mathbb{R}$ which evaluate the constancy of invariant image features at corresponding pixels in subsequent frames, e.g. the constancy of a) the brightness [14], b) the intensity, c) the gradient or d) the curvature at a given position $\boldsymbol{x} \in \Omega \times [0,T]$ with corresponding displacement vector $u \in \mathbb{R}^2$:

a) brightnessConst: $f(\boldsymbol{x},\boldsymbol{u}) = \nabla\mathcal{I}(\boldsymbol{x})\,\boldsymbol{u}(\boldsymbol{x}) + \frac{\partial \mathcal{I}(\boldsymbol{x})}{\partial t}$
b) ssdConst: $f(\boldsymbol{x},\boldsymbol{u}) = \|\mathcal{I}(\boldsymbol{x}) - \mathcal{I}_w(\boldsymbol{x})\|_{l_2}^2$
c) gradConst: $f(\boldsymbol{x},\boldsymbol{u}) = \|\nabla\mathcal{I}(\boldsymbol{x}) - \nabla\mathcal{I}_w(\boldsymbol{x})\|_{l_2}^2$
d) hessConst: $f(\boldsymbol{x},\boldsymbol{u}) = \|H(\boldsymbol{x}) - H_w(\boldsymbol{x})\|_{l_2}^2$

Here $\mathcal{I}_w$ and H_w denote the image sequence and the Hessian of the image sequence warped by the computed flow field $\boldsymbol{u}$. The set of invariance functions will be denoted by $\mathcal{E}$. A surface function for a given flow vector $\boldsymbol{u}$ reflects the variation of an invariance function $f \in \mathcal{E}$ over the set of modifications of the current displacement vector. Hence, it can be understood as an indicator for possible alternatives to the current displacement vector:

$$S_{\boldsymbol{x},\boldsymbol{u},f} : \mathbb{R}^2 \to [0,1],\; S_{\boldsymbol{x},\boldsymbol{u},f}(\boldsymbol{d}) := f(\boldsymbol{x}, \boldsymbol{u} + \boldsymbol{d}) \ . \tag{1}$$

The surface functions are used to derive surface measures for the detection of the id2-o situation based on any given invariance function $f \in \mathcal{E}$ and the following theoretical considerations. Low surface function values $S_{\boldsymbol{x},\boldsymbol{u},c}(\boldsymbol{d})$ for several displacements $\boldsymbol{d}$ denote uncertainty in the flow estimation, such as in the case of homogeneous regions and aperture problems. These situations can be detected by analyzing the curvature of the surface function along its principal axes, which is small along both axes for homogeneous regions and along one principal axis for aperture problems. In the case of occlusion, transparent structures and noise the minimum of the surface function is usually high indicating that no good estimate is possible at all. In contrast, a single, low minimum suggests a unique, reliable displacement vector. Hence, the intrinsic dimension of the surface function together with its minimum value yield information on the reliability of optical flow estimates. Only in the case of intrinsic dimension two and a low minimum value an optical flow estimate can be understood as reliable due to missing alternatives and a low value of the invariance function. The case of i1d and i0d can be detected by at least one low curvature value of the principal surface axes. The case of occlusion, severe noise and transparent structures yields a high minimum value of the surface function as none of the modified displacement vectors fulfills the assumption of the invariance function. Therefore, the confidence value should always be close to 1 if the smaller curvature value c_S is high and the surface minimum m_S is low. Let S be the set of surface functions defined in Equation (1). Then the surface measure can be defined as

$$m_f : \Omega \times \mathbb{R}^2 \to [0,1],\ m_f(\boldsymbol{x},\boldsymbol{u}) := \varphi(S_{\boldsymbol{x},\boldsymbol{u},f}) \tag{2}$$

$$\varphi : S \to [0,1],\ \varphi(S_{\boldsymbol{x},\boldsymbol{u},f}) := \frac{1}{1+m_S} \cdot (1 - \frac{1}{1+\tau c_S^2}) \ . \tag{3}$$

$\tau \in \mathbb{R}^+$ is used to scale the influence of the minimum curvature value. In this paper we use $\tau = 60$. The discretization of the surface measure has a large influence on the quality of the result. To discretize a surface function $S_{\boldsymbol{x},\boldsymbol{u},f}$ we use a step size h denoting the distance between two surface points and a fixed size b referring to the number of surface points in horizontal and vertical direction after discretization, e.g. $h = 0.5,\ b = 13$ yielded good results. For high accuracy h is chosen between 0 and 1 and bicubic interpolation is used. Examples for discretized surface functions are shown in Figure 1.

To obtain robustness towards noise we weight the surface by a Gaussian function centered on the central pixel before choosing the minimum value m_S and ignore all surface positions separated from the minimum by a local maximum for

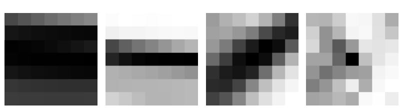

a) b) c) d)

Fig. 1. Discretized surface functions: a) i0d, b),c) i1d, d) i2d

the calculation of the principal surface axes. Since the eigenvalues of the Hessian yield noisy curvature estimates, a robust curvature estimator is introduced. It averages n curvature values along the principal axis using the following filter mask: $\frac{1}{n}(\underbrace{1\ldots1}_{n}\ -2n\ \underbrace{1\ldots1}_{n})$.

5 Motion Inpainting

We will show that sparsification based on the information contained in a surface measure map with subsequent motion inpainting improves the flow fields calculated with the combined local global and the structure tensor method on our test sequences (the well-known Marble and Yosemite sequence as well as the Street and Office sequence [15]). Motion inpainting is formulated as the minimization of a functional, which can be optimized by means of a variational approach. Let $\omega \subset \Omega \times [0,T]$ be a closed subset of the image domain, where the flow field is to be reconstructed. This is the case for all pixel positions which have been classified as unreliable by the surface measure. Let $\partial\omega$ be the boundary of ω, and let $\boldsymbol{u}^*$ be the reconstructed flow field within the region ω. Then $\boldsymbol{u}^*$ is the solution to the minimization problem

$$\min \int_{\omega} \|\nabla_3 \boldsymbol{u}^*\|^2 \text{ with } \boldsymbol{u}^*|_{\partial\omega} = \boldsymbol{u}|_{\partial\omega} \ . \tag{4}$$

Here ∇_3 means the spatio-temporal gradient. The minimizer satisfies the following system of Euler-Lagrange equations consisting of the Laplacian of each component of the flow field:

$$\triangle u_i = u_{i_{xx}} + u_{i_{yy}} + u_{i_{tt}} = 0,\ i \in \{1,2\},\ \boldsymbol{u}^*|_{\partial\omega} = \boldsymbol{u}|_{\partial\omega} \ . \tag{5}$$

We discretize the Euler-Lagrange equations using finite differences, von Neumann boundary conditions and the three-dimensional seven point Laplace stencil. The resulting linear system of equations is solved by successive overrelaxation.

6 Experiments and Results

To evaluate our results we first compare the quality of the surface measures to the previously known confidence measures detecting i2d situations and show that all of our surface measures - independent of the underlying invariance function - perform better than the best previously proposed measures and are robust to noise as well. Then we show motion inpainting results which demonstrate the usefulness of the suggested postprocessing method.

6.1 Comparison to i2d Measures

As test sequence we use the synthetic sequence on the left in Figure 2, as it contains every intrinsic dimension and occlusions of every intrinsic dimension.

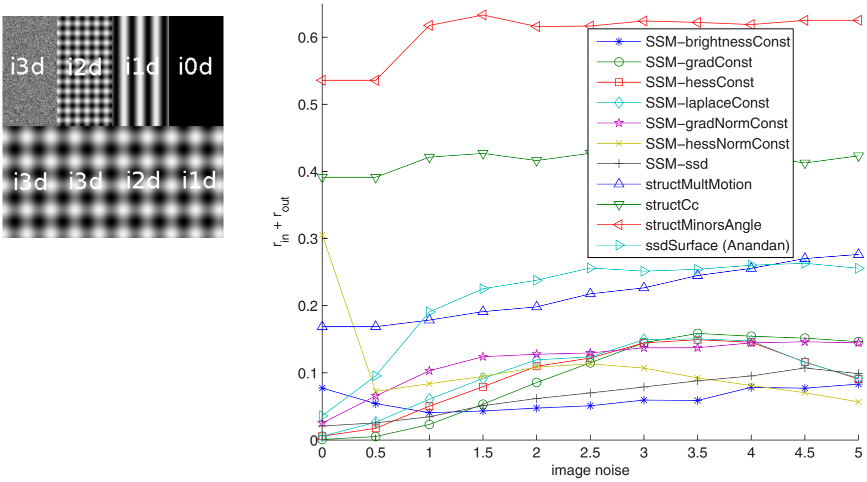

Fig. 2. Left: test sequence for surface measures containing four different intrinsic dimensions and the case of augmented intrinsic dimensions due to occlusion of the upper half in the lower half; Right: Comparison of surface measures (SSM-measures) based on different invariance functions for the recognition of i2d-o situations to known confidence measures for increasing noise levels

In this way we can examine if the surface measures are able to recognize the i2d-o situation and if they can distinguish it from occluded lower intrinsic dimensions. The patterns are moving to the lower right, and the lower half of the sequence is occluded by the sine pattern in the current frame in order to obtain examples of increased intrinsic dimensions due to occlusion.

To obtain numerical results let $g(\boldsymbol{x})$ denote the ground truth at position $\boldsymbol{x}$. We compute two values expressing the average deviation of the measure for the set of pixels P where the i2d-o situation is present and for the set of pixels Q where the situation is not present. The sum of both values serves as error measure.

$$r = r_{in} + r_{out} = \frac{\sum_{\boldsymbol{x}\in P} |m_c(\boldsymbol{x}, \boldsymbol{u}(\boldsymbol{x})) - g(\boldsymbol{x})|}{|P|} + \frac{\sum_{\boldsymbol{x}\in Q} |m_c(\boldsymbol{x}, \boldsymbol{u}(\boldsymbol{x})) - g(\boldsymbol{x})|}{|Q|} \quad (6)$$

As no measures are known for the detection of the i2d-o situation we compare our surface measures to the best known measures for the i2d situation: structMultMotion derived from [16], structCc [1], Anandan's measure [17] and structMinorsAngle [18]. Figure 2 (right) shows the error measure r plotted against an increasing noise level $\sigma \in [0, 5]$ in the test sequence. The proposed surface measures are labeled by the prefix "SSM" and an abbreviation of the invariance function they are based on. We can see that the proposed surface measures generally perform better than the best previously proposed i2d measures for any underlying invariance function f. All surface measures are robust to noise, but depend on the robustness of the underlying invariance function. The susceptibility to noise increases with the order of the derivatives in the invariance function.

However, the influence of noise on the surface measures is limited by the robust curvature estimation along the principal axes.

6.2 Application to Other Test Sequences

For further validation of the surface measures we apply them to standard test sequences. As no ground truth concerning the i2d-o situation is available for these sequences, only a visual evaluation is feasible.

Figure 3 a)-f) shows six different cropped regions of the Marble sequence and the corresponding surface measure result based on the brightness constancy invariance function (brightnessConst). In Figure 3 a), b) and c) we can see the application of the surface measure to different textures. In a) and b) the Marble blocks show only very little texture which makes these regions unreliable for flow estimation. In contrast, most parts of the block texture in c) are classified as sufficient for a reliable flow computation. In d) and e) we can see examples of aperture problems (i1d). The diagonal line on the table as well as the edges of the flagstones in the background of the sequence are typical examples for this situation. Both are recognized well by the surface measure. The corners of the flagstones are correctly recognized as i2d-o regions. The table region in f) is partially recognized as i2d-o and partially as i0d. This is due to the larger homogeneous regions in the table texture, as here the result depends on the size of the surface considered. If the whole surface function lies within the homogeneous region, the curvature along the main axis is 0 and thus the surface measure result as well. To demonstrate that our surface measures can also detect occlusions, we

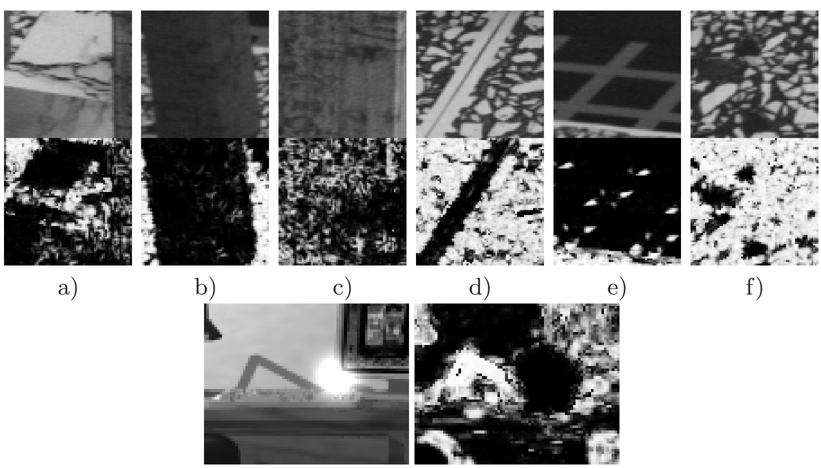

Fig. 3. Top: Cropped Marble sequence regions with result of brightness constancy surface measure for the recognition of i2d-o situations scaled to [0,1]; a),b),c) texture of blocks (i2d-o/i0d), d) diagonal table line (i1d) e) flagstones in the background (i1d,i2d-o at corners), f) table (i2d-o/i0d), Bottom: Office sequence with additional lens flare and result of the SSD constancy Surface Measure correctly idetifying the occlusion

use the cropped Office sequence [15] with an additional lens flare occluding part of the background in Figure 3 (this kind of lens flare often poses problems e.g. in traffic scenes). The brightness constancy surface measure detects this region.

6.3 Motion Inpainting

To evaluate the performance of the motion inpainting method we use a surface measure map to sparsify flow fields calculated on four ground truth sequences (Marble, Yosemite, Street and Office) by the three dimensional linear CLG method by Bruhn et al. [19] (a widely used global method) and by the structure tensor method by Bigün [20] (a fast local method). We apply motion inpainting to the sparsified displacement fields in order to reconstruct the flow at pixels with low surface measure values. We demonstrate that the angular error [4] is reduced significantly by means of motion inpainting. Table 1 shows the average angular error and standard deviation over ten frames for the sparsification and reconstruction of the flow field for the best previously proposed confidence measure (structMultMotion) and the new surface measures. For sparsification, we chose the flow field density optimal for motion inpainting with respect to the angular error.

Concerning the quality of the proposed measures, we can draw several conclusions from the results presented in Table 1. Firstly, the average angular error of the motion inpainting algorithm based on the surface measures is lower than the error we obtain based on the best previously proposed confidence measure. Hence, using the surface measures we can make more accurate statements on the reliability of the flow estimation than by means of previous i2d confidence measures. Secondly, the average angular error after motion inpainting is lower than the original angular error for the CLG and the structure tensor method. Thus, we conclude that the remaining flow vectors after sparsification contain all relevant information of the original flow field, and that any other information is dispensable, even obstructive, for the computation of a 100% dense flow field.

Table 1. Angular error for four test sequences for original field, sparsified field with given density, result of motion inpainting based on best surface measure and result of motion inpainting based on previously best confidence measure (structMultMotion), averaged over ten frames for the CLG and the structure tensor (ST) method; the density indicates the density optimal for motion inpainting

CLG	original	sparsification	density (%)	inpainting	previously best
Marble	3.88 ± 3.39	3.59 ± 3.03	70.6	3.87 ± 3.38	3.88 ± 3.39
Yosemite	4.13 ± 3.36	2.78 ± 2.24	20.7	3.85 ± 3.00	4.13 ± 3.36
Street	8.01 ± 15.47	2.77 ± 2.52	11.5	7.73 ± 16.23	7.99 ± 15.48
Office	3.74 ± 3.93	3.25 ± 4.80	26.7	3.59 ± 3.93	3.62 ± 3.91
ST	original	sparsification	density	inpainting	previously best
Marble	4.49 ± 6.49	2.96 ± 2.25	42.3	3.40 ± 3.56	3.88 ± 4.89
Yosemite	4.52 ± 10.10	2.90 ± 3.49	37.5	2.76 ± 3.94	4.23 ± 9.18
Street	5.97 ± 16.92	2.07 ± 5.61	34.6	4.95 ± 13.23	5.69 ± 16.47
Office	7.21 ± 11.82	2.59 ± 4.32	5.1	4.48 ± 4.49	6.35 ± 10.14

Finally, the table also indicates the average angular error for the sparsification of the flow field by means of the surface measures. Here we chose the sparsification density which has been found optimal for motion inpainting. The sparsification error is lower than the motion inpainting error and can be achieved if a dense flow field is not required. Hence, we have demonstrated that the quality of the surface measures is superior to previous measures and that the information contained in the remaining flow field after sparsification is sufficient for reconstruction.

Concerning the results of the motion inpainting algorithm we can draw the following conclusions. For both the CLG and the structure tensor method the sparsification of the flow field based on surface measures and subsequent inpainting yields lower angular errors than the original methods for all test sequences. The results of the local structure tensor method after motion inpainting are even superior to the original and the inpainted global CLG method in all cases but one. Therefore, we can conclude that - in contrast to the accepted opinion which favors global methods over local methods if dense flow fields are required - the filling-in effect of global methods is not necessarily beneficial for obtaining an accurate dense flow field. Instead, local and global methods alike can lead to better results if motion inpainting in combination with surface measures for sparsification is employed. Here, local methods often even seem preferable.

7 Summary and Conclusion

We have presented a method to estimate the feasibility of accurate optical flow computation. The proposed surface measures have proven robust to noise and are able to detect the situations where the full flow cannot be estimated reliably. They yield better results than previously proposed confidence measures and contain all relevant information for the reconstruction of the original flow field with even higher quality. Based on these measures we sparsified the original locally or globally computed flow field and filled in the missing flow vectors by a basic motion inpainting algorithm. Tests have been conducted using the CLG method and the structure tensor method on four standard test sequences. For our test sequences we can conclude that the application of a postprocessing method to sparsified flow fields calculated with local or global methods yields better results than can be achieved by exploiting the filling-in effect of global methods. Hence, in contrast to the accepted opinion, global methods are not always preferable to local methods if a dense flow field is required, because motion inpainting only based on reliable flow vectors can lead to superior results.

References

1. Haussecker, H., Spies, H.: Motion. In: Jähne, B., Haussecker, H., Geissler, P. (eds.) Handbook of Computer Vision and Applications, ch. 13, vol. 2. Academic Press, London (1999)
2. Bruhn, A., Weickert, J.: A Confidence Measure for Variational Optic Flow Methods, pp. 283–298. Springer, Heidelberg (2006)

3. Kondermann, C., Kondermann, D., Jähne, B., Garbe, C.: An adaptive confidence measure for optical flows based on linear subspace projections. In: Hamprecht, F.A., Schnörr, C., Jähne, B. (eds.) DAGM 2007. LNCS, vol. 4713, pp. 132–141. Springer, Heidelberg (2007)
4. Barron, J.L., Fleet, D.J., Beauchemin, S.: Performance of optical flow techniques. International Journal of Computer Vision 12(1), 43–77 (1994)
5. Bainbridge-Smith, R., Lane, A.: Measuring confidence in optical flow estimation. IEEE Electronics Letters 32(10), 882–884 (1996)
6. Rosenberg, A., Werman, M.: Representing local motion as a probability distribution matrix applied to object tracking. In: Proceedings of the International Conference on Computer Vision and Pattern Recognition, pp. 654–659 (1997)
7. Irani, M., Anandan, P.: Robust multi-sensor image alignment. In: Proceedings of the International Conference on Computer Vision, pp. 959–966 (1998)
8. Matsushita, Y., Ofek, E., Tang, X., Shum, H.: Full-frame video stabilization. In: Proceedings of the International Conference on Computer Vision and Pattern Recognition, vol. 1, pp. 50–57 (2005)
9. Bishop, C.: Neural Networks for Pattern Recognition. Oxford University Press, New York (1995)
10. Zetzsche, C., Barth, E.: Fundamental limits of linear filters in the visual processing of two dimensional signals. Vision Research 30(7), 1111–1117 (1990)
11. Barth, E.: Bewegung als intrinsische geometrie von bildfolgen. In: Proceedings of the German Association for Pattern Recognition (DAGM) (1999)
12. Barth, E., Stuke, I., Aach, T., Mota, C.: Spatio-temporal motion estimation for transparency and occlusions. In: Proceedings of the International Conference on Image Processing (ICIP), vol. 3, pp. 69–72 (2003)
13. Krüger, N., Felsberg, M.: A continuous formulation of intrinsic dimension. In: British Machine Vision Conference (2003)
14. Horn, B., Schunck, B.: Determining optical flow. Artificial Intelligence 17, 185–204 (1981)
15. McCane, B., Novins, K., Crannitch, D., Galvin, B.: On benchmarking optical flow (2001), `http://of-eval.sourceforge.net/`
16. Mota, C., Stuke, I., Barth, E.: Analytical solutions for multiple motions. In: Proceedings of the International Conference on Image Processing ICIP (2001)
17. Anandan, P.: A computational framework and an algorithm for the measurement of visual motion. International Journal of Computer Vision 2, 283–319 (1989)
18. Barth, E.: The minors of the structure tensor. In: Proceedings of the DAGM (2000)
19. Bruhn, A., Weickert, J., Schnörr, C.: Lucas/Kanade meets Horn/Schunck: Combining local and global optic flow methods. International Journal of Computer Vision 61(3), 211–231 (2005)
20. Bigün, J., Granlund, G.H., Wiklund, J.: Multidimensional orientation estimation with applications to texture analysis and optical flow. IEEE Journal of Pattern Analysis and Machine Intelligence 13(8), 775–790 (1991)