Study of Dynamical Processes with Tensor-Based Spatiotemporal Image Processing Techniques

B. Jähne[1,2,4], H. Haußecker[1], H. Scharr[1], H. Spies[1], D. Schmundt[1,3], and U. Schurr[3]

[1] Research Group Image Processing, Interdisciplinary Center for Scientific Computing, Heidelberg University, Im Neuenheimer Feld 368, 69120 Heidelberg
[2] Institute for Environmental Physics, Heidelberg University
Im Neuenheimer Feld 366, 69120 Heidelberg
[3] Institute of Botany, Heidelberg University
Im Neuenheimer Feld 360, 69120 Heidelberg,
[4] Scripps Institution of Oceanography, University of California, San Diego
La Jolla, CA, 92093-0230, USA,

Abstract. Image sequence processing techniques are used to study exchange, growth, and transport processes and to tackle key questions in environmental physics and biology. These applications require high accuracy for the estimation of the motion field since the most interesting parameters of the dynamical processes studied are contained in first-order derivatives of the motion field or in dynamical changes of the moving objects. Therefore the performance and optimization of low-level motion estimators is discussed. A tensor method tuned with carefully optimized derivative filters yields reliable and dense displacement vector fields (DVF) with an accuracy of up to a few hundredth pixels/frame for real-world images. The accuracy of the tensor method is verified with computer-generated sequences and a calibrated image sequence. With the improvements in accuracy the motion estimation is now rather limited by imperfections in the CCD sensors, especially the spatial nonuniformity in the responsivity. With a simple two-point calibration, these effects can efficiently be suppressed. The application of the techniques to the analysis of plant growth, to ocean surface microturbulence in IR image sequences, and to sediment transport is demonstrated.

1 Introduction

Since December 1995, the Interdisciplinary Center for Scientific Computing, the Institute for Environmental Physics, and the Institute for Botany cooperate in a DFG-funded interdisciplinary research unit ("Forschergruppe") to study transport, exchange, and growth processes. The oceanographic applications are investigated in close cooperation with the Scripps Institution of Oceanography, University of California, San Diego. The combination of novel visualization techniques and image sequence processing techniques gives an unprecedented insight into complex dynamic processes. This approach allows studying key scientific

questions for which previously no adequate experimental techniques were available. A close and interdisciplinary cooperation between applications and fundamental research in image analysis is the most distinct feature of the research unit. The application areas currently include

- small-scale air sea interaction processes, especially air-water gas transfer and wind-generated waves [14,11],
- plant leaf growth processes including the measurement of the physiological relevant parameters, and
- analysis of the pollution plumes (especially NO_x) from biomass burning and industrial areas using multispectral image sequences of the GOME instrument on the ERS2 satellite.

The "objects" encountered in the research unit differ from those normally studied in computer vision and thus pose new challenges for image sequence processing and analysis:

Accuracy. The displacements between consecutive frames must be determined with a relative accuracy much better than 1 %. This requirement results in an error for the interframe displacement of no larger than a few hundredth pixels.

Derivatives of motion field. For plant growth studies and the analysis of turbulent flows the optical flow is not of primary interest but spatial derivatives of the motion field such as divergence and vorticity.

Estimate of dynamical changes. The objects are non-rigid undergoing dynamical changes. New parts may gradually appear when studying the growth of plant leaves or roots or when new waves are generated by the wind at the ocean surface.

Multichannel motion estimation. An adequate study of dynamic processes often requires the simultaneous acquisition of image sequences of many parameters. Therefore the image sequence processing techniques must be set up in a way that they can also be used with multichannel images. This means more than just using multiple channels to determine a *single* velocity. Often the different channels show slightly *different* velocities, which are significant to characterize the underlying dynamical process. In the plumes from biomass burning, for example, different tracers have different life times. Or, if multiple tracers are used to study the transfer processes across the ocean surface, different diffusivities lead to different concentration patterns from which important clues about the underlying mechanisms can be drawn.

Speed of processing. Systematic studies of the dynamical processes investigated require processing of huge amounts of image sequences. Thus the algorithms used must also be fast.

The requirements summarized above caused us to revisit the basic approaches to image sequence processing. While there is a wealth of different concepts to compute optical flow (for a recent review see, e. g., [1]), much less work has been devoted to analyze the performance of optical flow algorithms with real-world image sequences and to optimize their implementation for accuracy and speed.

In this paper we report significant improvements in the accuracy of a tensor-based technique analyzing motion as orientation in spatiotemporal images. After briefly reviewing the tensor technique in Sect. 2, the focus in this paper is on an accurate and fast implementation in Sect. 3. One key point is a new nonlinear filter optimization technique minimizing the error in the direction of gradient operators (Sect. 3.2). In Sect. 4 the accuracy of the new approach is first verified with computer generated and calibrated real-world image sequences. Then it is shown that a careful radiometric calibration of the camera is required to avoid systematic errors in the motion field due to the photoresponse nonuniformity (PRNU) of the imaging sensor. Finally, various applications of the research unit are shown and discussed in Sect. 5.

2 Theory of the Structure Tensor

2.1 Motion as Orientation in Spatiotemporal Images

The displacement of gray value structures in spatiotemporal images yields inclined image structures with respect to the temporal axis. The relation between the orientation angle and the optical flow is given by

$$\boldsymbol{u} = -\left[\tan\varphi_1, \tan\varphi_2\right]^T, \tag{1}$$

where $\boldsymbol{u} = [u_x, u_y]^T$ denotes the optical flow on the image plane and the angles φ_1 and φ_2 define the angles between the plane normal to the lines of constant gray value and the x and y axes, respectively [7]. This basic property of spatiotemporal images allows to estimate the optical flow from a 3-D orientation analysis, searching for the orientation of constant gray value in $\boldsymbol{x}t$-space. Orientation is different from direction. While the direction (e. g., of a gradient vector) is defined for the full angular range of 360°, the orientation is invariant under a rotation of 180° and thus has only a definition range of 180°.

In order to determine local orientation locally, Bigün and Granlund [2] proposed a tensor representation of the local grey value distribution. Using directional derivatives, Kass and Witkin [12] came to a solution that turned out to be equivalent to the tensor method. Searching for a general description of local orientation Knutsson [13] concluded that local structure in an n-dimensional space can be represented by a symmetric $n \times n$ tensor. Rao [19] used a similar tensor representation for 2D texture analysis.

2.2 Total Least Squares Optimization

The orientation of gray-value structures can mathematically be formulated as a total least squares optimization problem [3]. The scalar product between a vector $\boldsymbol{r}$, representing the orientation for constant gray values in the image sequence, and the spatiotemporal gradient $\nabla_{xt} g(\boldsymbol{x}, t)$ is a semi-positive definite bilinear form that expresses the local deviation of the spatiotemporal gray value structure from an ideally oriented structure. If the gradient is perpendicular to

$\boldsymbol{r}$, the product is zero. It reaches a maximum when the gradient is either parallel or antiparallel to $\boldsymbol{r}$. Thus the expression

$$\int_{-\infty}^{\infty} h(\boldsymbol{x}-\boldsymbol{x}', t-t') \left(\boldsymbol{r}^T(\boldsymbol{x},t)\, \boldsymbol{\nabla}_{xt} g(\boldsymbol{x}', t')\right)^2 \mathrm{d}^2x'\mathrm{d}t', \tag{2}$$

has to be minimized in a local neighborhood. The size of the local neighborhood around the central point $(\boldsymbol{x},t)$ is given by the shape of the window-function $h(\boldsymbol{x}-\boldsymbol{x}', t-t')$. Equation (2) can be rewritten as a matrix equation

$$\boldsymbol{r}^T \left[h * \left(\boldsymbol{\nabla}_{xt} g\, \boldsymbol{\nabla}_{xt} g^T\right)\right] \boldsymbol{r} = \boldsymbol{r}^T \boldsymbol{J}\, \boldsymbol{r} \rightarrow \min, \tag{3}$$

with the 3D *structure tensor*

$$\boldsymbol{J}(\boldsymbol{x}) = \int_{-\infty}^{\infty} h(\boldsymbol{x}-\boldsymbol{x}', t-t')\, \boldsymbol{\nabla}_{xt} g(\boldsymbol{x}', t')\, \boldsymbol{\nabla}_{xt} g^T(\boldsymbol{x}', t')\, \mathrm{d}^2x'\mathrm{d}t'. \tag{4}$$

The components of this tensors are abbreviated by

$$\boldsymbol{G}_{pq} = h * \left(\frac{\partial g}{\partial p}\frac{\partial g}{\partial q}\right) \quad \text{with} \quad p,q \in \{x,y,t\} \tag{5}$$

Equation (3) reaches a minimum if the vector $\boldsymbol{r}$ is given by the *eigenvector* of the tensor $\boldsymbol{J}$ to the *minimum eigenvalue* [2,18]. The estimation of local orientation therefore reduces to an *eigenvalue analysis* of the tensor $\boldsymbol{J}$.

The eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3$ of the tensor can be used to characterize the spatiotemporal gray value structure [6], where the smallest eigenvalue λ_3 reflects the noise level in the image sequence. If it is significantly higher than expected from the noise level, the neighborhood shows no constant motion. If the smallest eigenvalue is consistent with the noise level and the two others are significantly larger than the noise level, both components of the optical flow can be estimated and are given by the eigenvector $\boldsymbol{e}_s$ to the smallest eigenvalue:

$$\boldsymbol{u} = \left(\frac{e_{s,x}}{e_{s,t}}, \frac{e_{s,y}}{e_{s,t}}\right). \tag{6}$$

If two eigenvalues are in the order of the noise level, an image structure with linear symmetry (spatial local orientation) moves with a constant velocity. This is the well known *aperture problem* in optical flow computation. The eigenvector to the largest eigenvalue, $\boldsymbol{e}_l$, points normal to the plane of constant gray values in the spatiotemporal domain and can be used to compute the normal optical flow $u_\perp$:

$$u_\perp = -\frac{e_{l,t}}{\sqrt{e_{l,x}^2 + e_{l,y}^2}}. \tag{7}$$

The tensor approach accounts for errors in *all* components of the spatiotemporal gradient. In contrast, the classical differential approach [15]

$$\begin{bmatrix} G_{xx} & G_{xy} \\ G_{xy} & G_{yy} \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \end{bmatrix} = - \begin{bmatrix} G_{xt} \\ G_{yt} \end{bmatrix} \tag{8}$$

which is a solution of the least-squares minimization problem

$$h * \left(\boldsymbol{u} \nabla g + \frac{\partial g}{\partial t} \right)^2 \rightarrow \min, \tag{9}$$

accounts only for errors in the temporal derivate $\partial g / \partial t$.

2.3 Analytical Performance Analysis

It was shown analytically by [7,9] that the tensor method — as well as other approaches — yield *exact* estimates of the optical flow under the ideal condition of a spatial pattern moving with a constant velocity. The different techniques, however, distinguish themselves in the way they respond to deviations and distortions from the ideal case. The tensor method shows — in contrast to the standard least squares differential approach — no bias of the estimate in images with isotropic normal distributed noise. Inhomogeneous and accelerated motion also causes a surprisingly low bias in the motion estimate [9].

The analysis summarized here refers to continuous spatiotemporal images. It thus does not include errors caused by the discretization. In Sect. 3.2 we will discuss that the largest error in the motion estimate is due to an inadequate discretization of the spatial and temporal derivates.

3 Accurate and Fast Implementation

3.1 Computing the Structure Tensor by Spatiotemporal Filtering

The implementation of the tensor components can be carried out very efficiently by simple spatiotemporal filter operators. Identifying the convolution in (4) with a 3D spatiotemporal smoothing of the product of partial derivatives with the window function $h(\boldsymbol{x}, t)$, each component of the structure tensor can be computed as

$$J_{pq} = \mathcal{B} \left(\mathcal{D}_p \cdot \mathcal{D}_q \right), \tag{10}$$

with the 3D spatiotemporal smoothing operator $\mathcal{B}$ and the differential operators $\mathcal{D}_p$ in the direction of the coordinate x_p.

Using a binomial operator the smoothing can be performed very efficiently on a multigrid data structure [5,8]. With this approach, the displacement vector field is stored on the next coarser level of a pyramid. Thus the number of computations for any subsequent processing — especially the eigenvalue analysis (Sect. 3.3) — is also reduced by a factor of four.

The smoothing and derivative filters can be computed very efficiently in integer arithmetic using modern multimedia instruction sets such as Intel's MMX or Sun's VIS [10].

3.2 Optimal Derivative Filters

The most critical point is the choice of an appropriate differential operator. A difference operator is only an approximation of a derivative operator. For an accurate computation of the structure tensor only deviations in the direction of the gradient are of importance as can be seen directly from (2). This means that for an optimal difference filter, it is not required to approximate the ideal transfer function of the derivative filter, ik_q. It suffices to approximate a more general transfer function

$$\hat{D}_q(\boldsymbol{k}) = \hat{f}(|\boldsymbol{k}|)k_q, \tag{11}$$

where $\hat{f}(|\boldsymbol{k}|)$ is an arbitrary isotropic function and k_q is the component of the wave number in the direction q (for the time axis this is the circular frequency ω). The transfer function (11) includes the classes of all regularized derivative filters that smooth the images before the discrete differences are computed.

The standard symmetric difference filter is a bad choice to compute the structure tensor. This filter has the transfer function

$$\hat{D}_q = \mathrm{i}\sin(\pi\tilde{k}_q), \tag{12}$$

where $\tilde{k}$ is the wave number normalized to the Nyquist wave number. Consequently the direction of the 2-D spatial gradient is given by

$$\phi = \arctan\frac{\sin(\pi\tilde{k}\sin\phi)}{\sin(\pi\tilde{k}\cos\phi)}. \tag{13}$$

A Taylor expansion of this expression in $\tilde{k}$ gives the angle error, $\Delta\phi$, in the approximation of small $\tilde{k}$ as

$$\Delta\phi = \frac{(\pi\tilde{k})^2}{24}\sin 4\phi + O(\tilde{k}^4). \tag{14}$$

The error is substantial, it is more than 5° at $\tilde{k} = 0.5$. The error of other standard derivative filters is not much lower. The error for the Sobel filter, for example, is only two times lower in the approximation for small wave numbers [8].

The general form of the transfer function in (11) suggests that an ansatz of a derivative filter with cross-smoothing in all directions except the direction in which the derivation is applied seems to be a good choice. (The symmetric difference filter contains already a certain degree of smoothing.) For the sake of an efficient computation, separable filters of the form:

$$\boldsymbol{D}_{x,\mathrm{opt}} = \boldsymbol{D}_x * \boldsymbol{B}_y * \boldsymbol{B}_t, \quad \boldsymbol{D}_x = \frac{1}{2}[1,0,-1], \quad \boldsymbol{B}_{y,t} = [p/2, 1-p, p/2]^T \tag{15}$$

are used with the transfer function

$$\hat{D}_{x,\mathrm{opt}}(p) = \mathrm{i}\sin(\pi\tilde{k}_x)((1-p) + p\cos(\pi\tilde{k}_y))\,((1-p) + p\cos(\pi\tilde{\omega})) \tag{16}$$

The equations above show the separable filter for derivation in x direction. The derivative filters in the other direction just show permuted indices for the directions.

The ansatz in (15) and (16) has one degree of freedom. It is optimized with a nonlinear functional that describes the deviation of the direction of the gradient built with the filters and contains an arbitrary wave number dependent weighting function so that the filters can be optimized for selected wave number ranges. Further details of the optimization technique and its generalization to larger filter kernels can be found in [20]. With rational filter coefficients, a value of $p = 6/16$ turned out to be an optimal choice. For wave numbers $\tilde{k}$ up to 0.5, the error in the direction of the gradient is well below 0.4°. The maximum error is thus more than ten times lower than for the Sobel filter. The nonlinear optimization procedure used here also produces significantly better filters than a similar approach by Farid and Simoncelli [4] based on a linear optimization technique.

3.3 Fast Eigenvalue Analysis

The *Jacobi transformation* [17] proved to be an efficient and robust algorithm to determine the eigenvalues and eigenvectors of the structure tensor. In all cases no more than 8 elementary Jacobi rotations were required. Interestingly, the Jacobi algorithm required a maximum of only three iterations with a spatially oriented gray value structure (aperture problem, rank one structure tensor). Since edges are generally more common than corners, the Jacobi algorithm is thus very efficient in the mean.

Preselecting interesting image regions could further speed up the eigenvalue analysis. If the *trace* of the structure tensor (sum of the eigenvalues) is below a certain noise dependent threshold, it indicates a homogeneous region in which the computation of the eigenvalues and eigenvectors is not required since no significant motion estimate can be retrieved then.

4 Results

4.1 Computer Generated Image Sequences

First computer generated sequences were used to test the ultimate accuracy of the implementation of the tensor method. In order to get realistic spatial structures, we used frame 13 (Fig. 2a) of the calibrated sequence taken by the IITB in Karlsruhe [16]. From this image a sequence was produced using optimized interpolation filters [8]. The result for a motion in x direction with 0.456 pixels/frame of a sequence to which zero mean normal distributed noise with a standard deviation of 2.0 bits was added is shown in Fig. 1. For the optimized filter and the Sobel filter, the standard deviation of the error in the velocity estimate is well below 0.01 pixels/frame. In contrast the simple symmetric filter shows both a systematic offset of almost 0.1 pixels/frame and a significant wider distribution in x direction. This test clearly shows that highly accurate motion estimation is possible under ideal condition with a constant motion field.

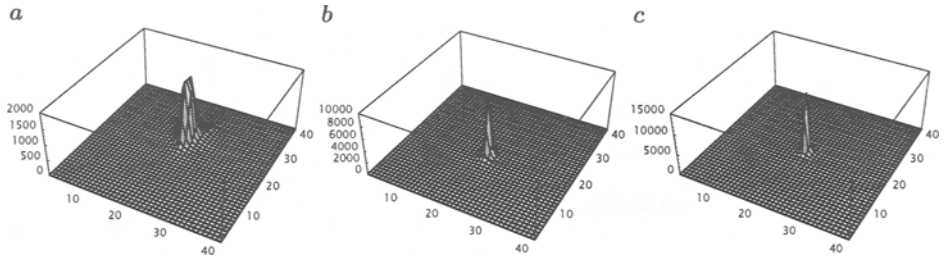

Fig. 1. 2-D histogram of the error in the motion estimation for a constant shift of 0.456 pixels/frame of one image of the IITB sequence. In addition, a zero-mean normal distributed noise with a standard deviation of 2.0 was added to each image of the computer generated sequence. One unit in the diagram corresponds to 0.01 pixels/frame, zero error is in the center at [20, 20]. Results with the tensor technique using **a** a symmetric derivative filter 1/2[1 0 -1], **b** a 3-D Sobel filter, and **c** the optimized $3 \times 3 \times 3$ filter described in Sect. 3.2.

4.2 Calibrated Real-world Sequence

The IITB image sequence [16] is one of the few calibrated image sequences available. It has also been used to test the error in the motion estimation with real image sequences. Fig. 2 shows the results. At the highly textured floor, a dense velocity field could be computed, while it is rather sparse at the marble columns (Fig. 2c). The smallest eigenvalue of the structure tensor shows the deviation from an ideally spatiotemporal structure with constant motion. Therefore parts of the edges of the columns appear bright in Fig. 2d. The error maps for the velocity estimates are shown in Fig. 2e and f, respectively. It is not surprising that the errors are high at motion discontinuities, because the structure tensor averages over a local spatiotemporal neighborhood defined by the size of the smoothing mask in (10). Again the much better performance of the optimized filter can be observed. The simple symmetric difference filter (Fig. 2e) shows a considerable systematic error that grows with increasing shifts in x direction (see Fig. 2b).

A 2-D histogram of the error maps for the velocity estimate is shown for three different derivative filters in Fig. 3. The optimized filter results in a standard deviation less than 0.03 pixels/frame. This is about a two times lower standard deviation than the best results in Table 1 from Otte and Nagel [16]. Their results are based on averaging over a $5 \times 5 \times 5$ neighborhood, while the smoothing used with the multigrid binomial filters in the tensor method effectively averages only over a sphere with a radius of 1.73 pixels or a volume of 22 pixels and thus preserves a higher spatial resolution. The non-optimized filters show significantly broadened distributions with standard deviations comparable to those found by Otte and Nagel [16] (Fig. 3a, b).

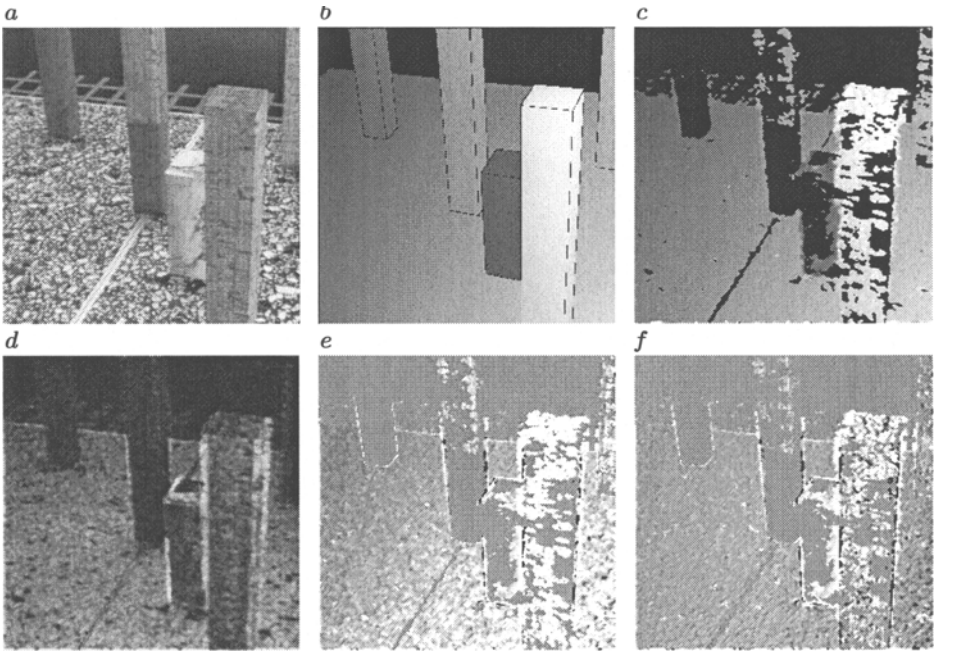

Fig. 2. Motion analysis using the calibrated sequence from the IITB, Karlsruhe [16]. **a** Frame 13 of the sequence. **b** x component of the true motion field. **c** Motion field determined with the tensor technique. Areas where no 2-D velocity could be determined are masked out. **d** Smallest eigenvalue of the structure tensor. **e** Error map (range [-0.2, 0.2] pixels/frame) for the 1/2 [1 0 -1] filter. **f** same as **e** for the optimized $3 \times 3 \times 3$ derivative filter (Sect. 3.2).

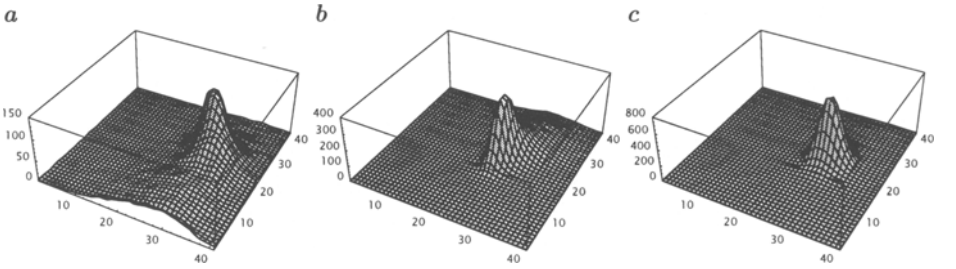

Fig. 3. 2-D histogram of the error in the motion estimation for the calibrated IITB sequence. Results with the tensor technique using **a** a symmetric derivative filter 1/2[1 0 -1],. **b** a 3-D Sobel filter, and **c** the optimized $3 \times 3 \times 3$ filter.

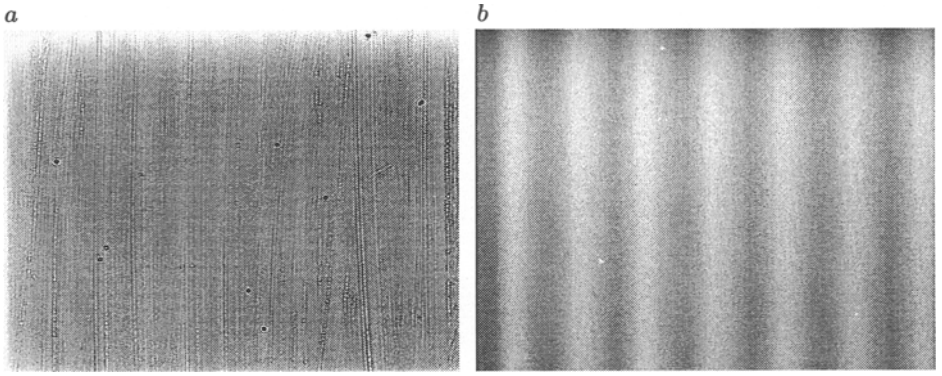

Fig. 4. Radiometric calibration of a standard CCD camera (Jai CV-M10): **a** Image of the relative responsivity in a range of ±3%. **b** Fixed pattern noise (without any illumination) in a range of ±0.5.

4.3 Influence of Sensor Imperfections

One may wonder why the standard deviations with real sensor data are significantly higher than those obtained with computer generated sequences (compare Fig. 1 und 3). The higher errors are related to imperfections of the CCD sensor/camera system.

A radiometric calibration study showed that standard CCD cameras show significant large-scale and small-scale spatial variations in the order of about 1%, which cannot be neglected (Fig. 4a). The fixed pattern noise without illumination is less serious. In the example shown in Fig. 4b the amplitude of the patterns is well below one. Since these patterns are static, they are superimposed to the real motion in the sequence. In parts of the sequence where the local contrast is low, the static patterns dominate the structure and thus a lower or even zero velocity is measured.

The influence of static patterns can nicely be demonstrated if objects with low contrast are moving such as the slightly textured elephant in Fig. 5. On the glass window of the CCD sensor, dirt causes spatial variations in the responsivity of the sensor (Fig. 5a, d). At the edges of the speckles, the smallest eigenvalue of the structure tensor shows high values indicating motion discontinuities (Fig. 5b). The motion field indeed shows drops at the positions of the speckles (Fig. 5c). If a simple two-point calibration is performed using the measured responsivity and an image with a dark pattern, the influence of the speckles is no longer visible both in the smallest eigenvalue and the motion field (Fig. 5e, f)

5 Application Examples

The structure tensor technique was applied to a variety of application examples from oceanography (IR ocean surface images), botany (growth processes), and

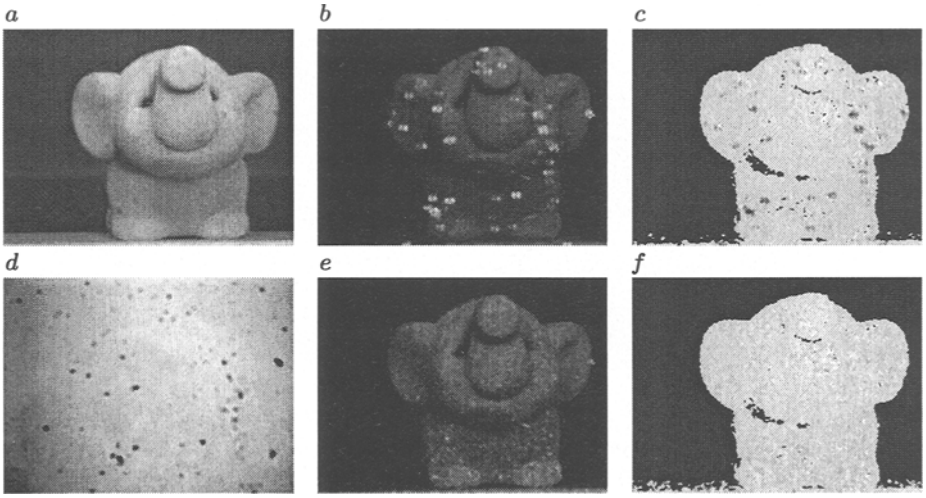

Fig. 5. Demonstration of the influence of spatial sensitivity variations of the CCD sensor on motion estimation: **a** One image of the elephant sequence. **b** Lowest eigenvalue of the structure tensor. **c** Velocity component in x direction. **d** Contrast enhanced relative responsivity. **e** Lowest eigenvalue of the structure tensor for a sequence corrected for the spatial responsivity changes. **f** Velocity component in x direction for the corrected sequence.

to traffic scenes. It proved to work well without any adaptation to the image content. Fig. 6 shows example motion fields of such sequences. The application to the IR sequences (Fig. 6a) demonstrates that the tensor technique is also suitable for noisy imagery.

The accuracy of the estimated velocity can also be demonstrated by computing first-order derivatives of the velocities. These are the most interesting quantities for the study of the dynamical processes we are investigating. The divergence of the motion field, for instance, directly gives the area-resolved growth rate (relative increase in area, see Fig. 7a). For the study of the microscale turbulence at the ocean surface, the divergence of the motion field is related to divergence and convergence zones at the interface (Fig. 7a).

Another interesting example is the study of sediment transport in the beds of rivers. Figure 8 shows the evaluation of a sequence where a strong flow in the river induced the transport of sand particles in the sediment. Despite the bad quality of the sequence taken by an endoscope put into the sand, reliable motion fields could be computed. The nature of these processes becomes evident in the images showing the divergence and rotation of the flow.

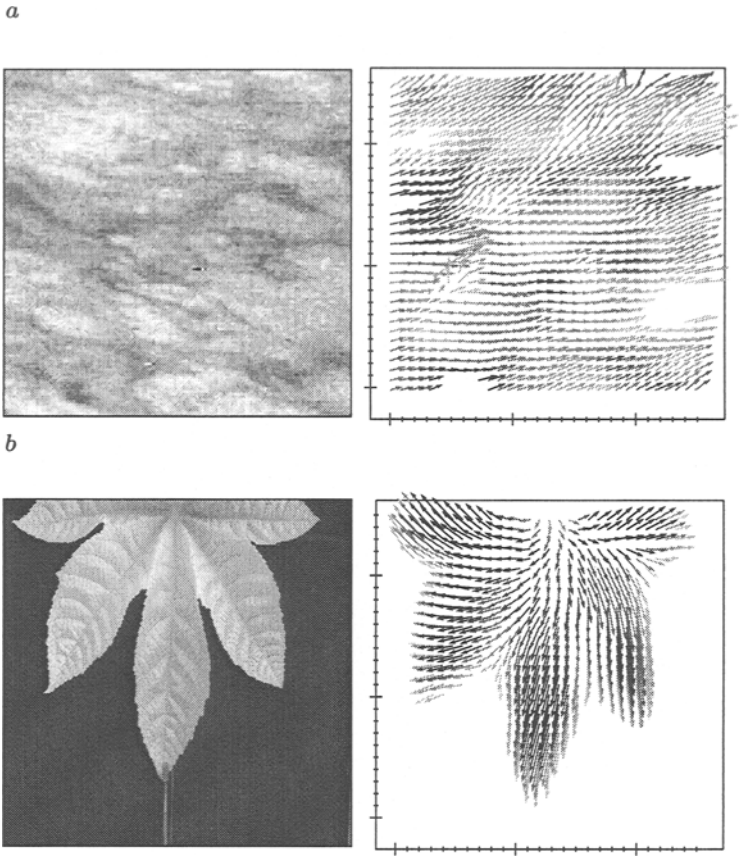

Fig. 6. Displacement vector fields computed with the tensor method. **a** IR images of the ocean surface for heat transfer studies, **b** growing leaf of a castor-oil plant. The arrows visualize the computed displacement field. Black arrows mean high certainty, lighter arrows lower certainty.

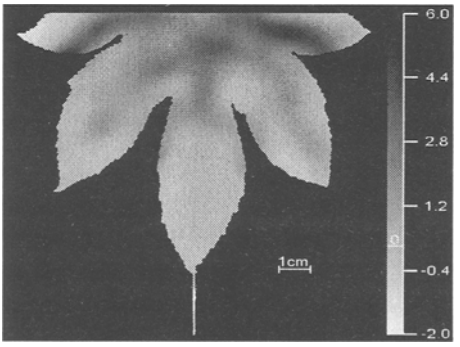

Fig. 7. Divergence of the motion field of a growing leaf of a castor-oil plant. The scale for the divergence ranges from -2.0 to 6.0 permille/min.

6 Conclusions

In this paper some recent research results of an interdisciplinary research unit were discussed. It is shown that the accuracy of motion estimation can significantly be improved over previous approaches by using a tensor method with carefully optimized derivative filters. The tensor approach also gives direct measures for the quality of the estimated motion. Some of the next steps of our research will include the extension to model-based motion estimation, multichannel image sequence processing, and the direct modeling of parameters of the dynamical processes. It is the long-term goal of our research to merge image sequence processing, 3-D image processing, and spectroscopic imaging into 5-D imaging as a powerful tool to study complex scientific problems.

Acknowledgments

Financial support for this research by the German Science Foundation (Deutsche Forschungsgemeinschaft, DFG, Jae 395/6-1) through the research unit (Forschergruppe) "Bildfolgenanalyse zum Studium dynamischer Prozesse", the US National Science Foundation (Coastal Ocean Processes, CoOP) and the US Office of Naval Research (ONR) through the Marine Boundary Layer Accelerated Research Initiative (MBL/ARI) is gratefully acknowledged. Sediment transport is being studied in a project funded by the German Federal Waterways Engineering and Research Institute (BAW), Karlsruhe (Dipl.-Ing. H.-J. Köhler). H. S. is supported by a scholarship within the DFG-Graduiertenkolleg "Modellierung und Wissenschaftliches Rechnen in Mathematik und Naturwissenschaften".

References

1. S. S. Beauchemin and J. L. Barron, 1997. The computation of optical flow, *ACM Computing Surveys, 27(3)*, pp. 433–467.
2. J. Bigün and G. Granlund, 1987. Optimal orientation detection of linear symmetry. Proc. *First Intern. Conf. on Comp. Vision, ICCV'87*, London, June 8-11, 1987, pp. 433–438.
3. C. H. Chu and E. J. Delp, 1989. Estimating displacement vectors from an image sequence, *J. Opt. Soc. Am. A6(6)*, pp. 871–878.
4. H. Farid and E.P. Simoncelli, 1997. Optimally rotation-equivariant directional derivative kernels. *7th Int. Conf. Computer Analysis of Images and Patterns*, Kiel, pp. 207–214.
5. H. Haußecker, 1995. Mehrgitter-Bewegungssegmentierung in Bildfolgen mit Anwendung zur Detektion von Sedimentverlagerungen, Diploma thesis, Univ. Heidelberg.
6. H. Haußecker and B. Jähne, 1997. A tensor approach for precise computation of dense displacement vector fields, Proc. *Mustererkennung 1997*, Braunschweig, 15–17. September 1997, E. Paulus und F. M. Wahl (Hrsg.), Informatik Aktuell, Springer, Berlin, pp. 199–208.
7. B. Jähne, 1993. *Spatio-Temporal Image Processing, Theory and Scientific Applications*, Lecture Notes in Computer Science, Vol. 751, Springer, Berlin, 1993.

8. B. Jähne, 1997. *Practical Handbook on Digital Image Processing for Scientific Applications*, CRC-Press, Boca Raton, FL, USA.
9. B. Jähne, 1997. Performance characteristics of low-level motion estimators in spatiotemporal images, Proc. *DAGM-Workshop Performance Characteristics and Quality of Computer Vision Algorithms*, Braunschweig, September 18, 1997, W. Foerstner (ed.), Institute of Photogrammetry, Univ. Bonn.
10. B. Jähne, 1997. SIMD-Bildverarbeitungsalgorithmen mit dem Multimedia Extension-Instruktionssatz (MMX) von Intel, *Automatisierungstechnik AT, 10*, pp. 453–460.
11. B. Jähne and H. Haußecker, 1998. Air-Water Gas Exchange, *Annual Rev. Fluid Mech., 30*, pp. 443–468.
12. M. Kass, and A. Witkin, 1987. Analyzing Oriented Patterns, *Comp. Vision Graphics Image Proc., 37*, pp. 362–385.
13. H. Knutsson, 1989. Representing local structure using tensors, *6th Scandinavian Conf. Image Analysis*, Oulu, Finland, pp. 244–251.
14. P. S. Liss and R. A. Duce (eds.). *The Sea Surface and Global Change*, Cambridge Univ. Press, Cambridge, UK, 1997.
15. B. D. Lucas, and T. Kanade. An iterative image-registration technique with an application to stereo vision, Proc. *DARPA Image Understanding Workshop*, pp. 121–130.
16. M. Otte, and H.-H. Nagel, 1994. Optical flow estimation: advances and comparisons, Proc. *ECCV'94, Vol. II*, J. O. Eklundh (ed.), Springer, Berlin, pp. 51–60.
17. W. H. Press, S. A. Teukolsky, W. T. Vetterling, B. P.:Flannery, 1992. *Numerical recipes in C: The Art of Scientific Computing*, Cambridge Univ. Press
18. A. R. Rao and B. G. Schunck, 1989. Computing oriented texture fields, *Proceedings CVPR'89, San Diego, CA*, pp. 61–68, IEEE Computer Society Press, Los Alamitos.
19. A. R. Rao, 1990. *A Taxonomy for Texture Description and Identification*, Springer, New York
20. H. Scharr, S. Körkel, and B. Jähne, 1997. Numerische Isotropieoptimierung von FIR-Filtern mittels Querglättung, Proc. *Mustererkennung 1997*, Braunschweig, 15–17. September 1997, E. Paulus und F. M. Wahl (Hrsg.), Informatik Aktuell, Springer, Berlin, pp. 367–374.

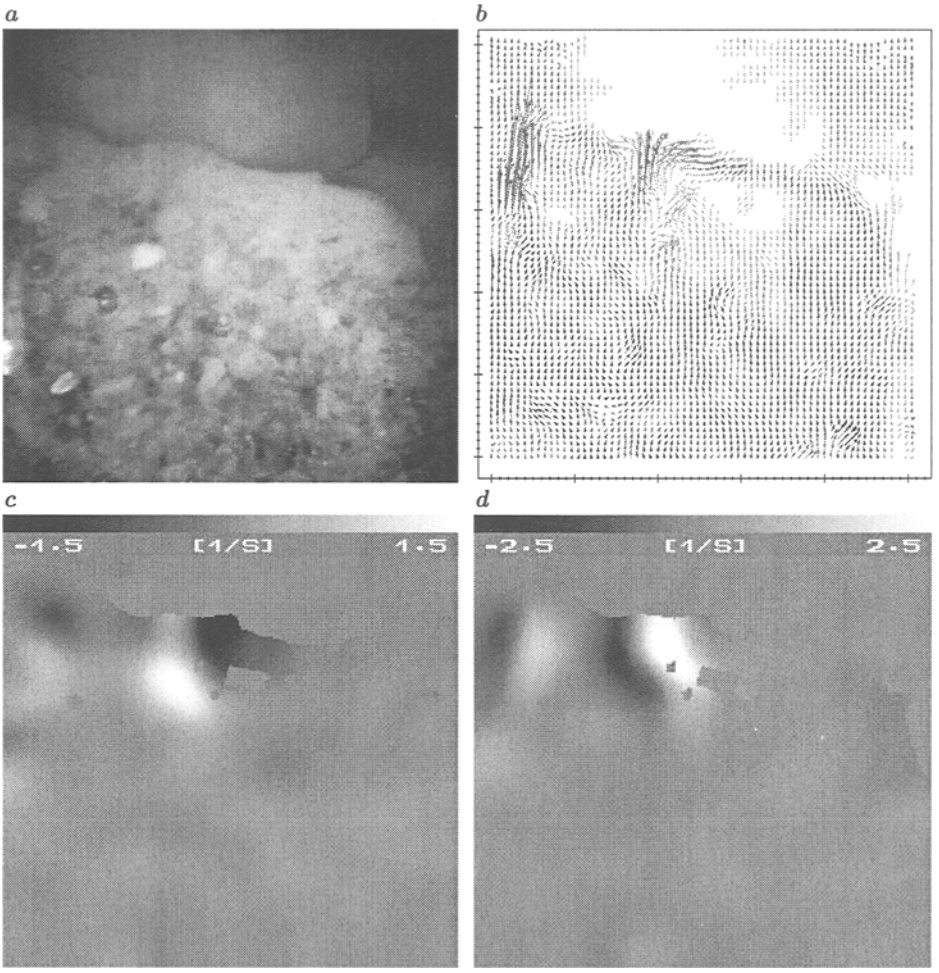

Fig. 8. Example for the study of flow in sediments. **a** One of the images of the flow of sand particles in the sediment observed with an embedded endoscope. **b** Displacement vector field computed with the tensor method. **c** Divergence and **d** rotation of the vector field as a color overlay on the original image; a color scale is included at the top of each image.