A Variational Approach to Adaptive Correlation for Motion Estimation in Particle Image Velocimetry[*]

Florian Becker[1], Bernhard Wieneke[2], Jing Yuan[1], Christoph Schnörr[1]

[1] Image and Pattern Analysis Group, Heidelberg Collaboratory for Image
Processing, University of Heidelberg, Germany,
{becker,yuanjing,schnoerr}@math.uni-heidelberg.de
[2] LaVision GmbH, Göttingen, Germany, bwieneke@lavision.de

Abstract. In particle image velocimetry (PIV) a temporally separated image pair of a gas or liquid seeded with small particles is recorded and analysed in order to measure fluid flows therein. We investigate a variational approach to cross-correlation, a robust and well-established method to determine displacement vectors from the image data. A "soft" Gaussian window function replaces the usual rectangular correlation frame. We propose a criterion to adapt the window size and shape that directly formulates the goal to minimise the displacement estimation error. In order to measure motion and adapt the window shapes at the same time we combine both sub-problems into a bi-level optimisation problem and solve it via continuous multiscale methods. Experiments with synthetic and real PIV data demonstrate the ability of our approach to solve the formulated problem. Moreover window adaptation yields significantly improved results.

1 Introduction

Overview. Particle image velocimetry is an important measurement technique for industrial fluid flow questions. Small particles are introduced into liquids or gases and act as indicators for the movement of the investigated substance around obstacles and in mixing zones. A 2D plane is illuminated by laser light rendering the particles in there visible to a camera which records two images of the highlighted area within a short time interval.

The analysis of the image data allows to determine the movement of particles and with this to measure the speed, turbulence or other derived mechanical properties of the fluid. In contrast to particle tracking velocimetry where first single objects are identified by their position and then matched between two image frames, algorithms for PIV determine patches from the first and second frame that fit best to some similarity measure. Cross-correlation has developed the state-of-the-art method for motion estimation in PIV and benefits from its robustness against noise and illumination disturbances.

In this paper we describe a variational approach to cross-correlation by continuously optimising over the displacement variables. In addition a criterion is defined to locally adapt the correlation window in order to improve accuracy of the estimation.

Related Work and Contribution. A vast number of literature exists on all aspects of the application of cross-correlation for PIV, here we only refer to [1] for an excellent overview. Typically an exhaustive search over the integer displacements is performed to search for the highest correlation peak. The correlation function is interpolated to gain sub-pixel accuracy. In contrast we present a variational approach to motion estimation based on *continuously* maximising the cross-correlation between two images. The correlation window is described by a "soft" Gaussian weighting function instead of a sharp rectangular mask. This idea is used both in a local [2] and global context [3] for smoothing the optical flow constraint. However while most approaches use a fixed window size common for all positions, we formulate a sound criterion for the location-dependent choice of the window shape parameters (size, orientation, anisotropy) in words of a further optimisation problem. Both displacements and window parameters are determined as a solution of a *combined*, bi-level minimisation problem which is being solved via a multiscale gradient-based algorithm. We test our approach with synthetic and real particle images to demonstrate the ability to robustly determine displacements and that window shape adaptation can improve results significantly.

Organisation. In section 2 we explicitly define the cross-correlation in words of a continuous minimisation problem and introduce the utilised weighting function. Our approach to window adaptation is motivated and described. Discretisation and optimisation of the defined approach is subject to section 3. Results of experiments with real and synthetic data are given in section 4. We conclude and describe further work in section 5.

2 Problem Statement

2.1 Variational Approach to Correlation

The input data consists of two images defined on $\Omega \subset \mathbb{R}^2$. However we define them to vanish outside Ω and thus obtain two infinite image functions $g_1, g_2 : \mathbb{R}^2 \mapsto \mathbb{R}$. For the continuous case, we define the negative cross-correlation function at position $x \in \Omega$ by

$$C(v, \Sigma, x) := -\int_{\mathbb{R}^2} w(y - x, \Sigma)\, g_1\left(y - \frac{1}{2}v\right) g_2\left(y + \frac{1}{2}v\right) \mathrm{d}y, \qquad (1)$$

with a window function $w(x, \Sigma)$ parametrised by Σ. In order to estimate the movement between two image frames in the considered areas, the correlation

function is minimised with respect to the displacement vector $v \in \mathbb{R}^2$. The local estimation is extended to a global variational problem to determine a vector field $u : \Omega \mapsto \mathbb{R}^2$,

$$\min_u \ C(u, \Sigma), \quad \text{with } C(u, \Sigma) := \int_\Omega C(u(x), \Sigma(x), x) \, \mathrm{d}x \,, \tag{2}$$

where Σ is defined on Ω and describes the location-dependent window shape. This formulation allows to add regularisation terms such as physical priors, e.g. incompressibility constraints [4], on the vector field. However here the integration over the correlation window in (1) is the only spatial regularising mechanism.

In this work we jointly determine the integer and fractional part of the displacement by *continuously* searching for an optimum of the correlation function. In addition we choose a "soft" window function $w(x, \Sigma) := \exp\left(-\frac{1}{2}x^\top \Sigma^{-1} x\right)$ which is basically a non-normalised Gaussian function, instead of a sharp, rectangular window. The symmetric, positive definite two-by-two matrix $\Sigma \in \mathcal{S}_{++}^2$ allows to continuously steer the size, anisotropy and orientation of the window, see Fig. 1 for some possible shapes.

Fig. 1. Possible shapes of the weighting function: varying size, anisotropy, orientation

2.2 Window Adaptation

When cross-correlation is employed for estimating motion it is implicitly assumed that the displacements within the considered window are homogeneous. However this only holds true in very simple cases and leads to estimation errors in areas of large motion gradients as the vector field is smoothed out. This effect could be avoided by reducing the window size, however with the harm of a smaller area of support and number of particles and thus a higher influence of image noise.

In order to improve accuracy we propose to adapt the window shape by minimising a function which models the expected error subject to the choice of the window parameter Σ at position $x \in \Omega$. Given a fixed vector field u we define the energy function as

$$E(\Sigma, u, x) := \int_{\mathbb{R}^2} w(y - x, \Sigma)e(x, y, u) \, \mathrm{d}y + \frac{\sigma^2}{2\pi\sqrt{\det \Sigma}} \,, \tag{3}$$

$$e(x, y, u) := \begin{cases} \|u(x) - u(y)\|_2^2 & \text{if } y \in \Omega \\ e_{\text{out}} & \text{otherwise} \end{cases}.$$

The first term of (3) measures the deviation from the assumption $u = \text{const}$. In addition we assume a constant error e_{out} if the correlation window incorporates data from outside the image domain.

The second term describes the error caused by insufficient large support for the displacement estimation in the presence of a homogeneous movement. We assume that each vector u results from a weighted least-square estimation $u = \arg\min_u \int w(x, \Sigma) \|u - \overline{u}(x)\|_2^2 \, dx$ over independent measurements $\overline{u}$ of the true displacement u^*, which are disturbed by Gaussian additive noise, i.e. $\overline{u}(x) \sim \mathcal{N}\left(u^*, \sigma^2 I\right)$. Then it is possible to show that the expected square error is $\mathcal{E}\left\{\|u - u^*\|_2^2\right\} = \frac{\sigma^2}{\int w(x, \Sigma) dx} = \frac{\sigma^2}{2\pi\sqrt{\det \Sigma}}$. The parameter σ in E constitutes an estimation of the influence of image noise on the measurement error.

Note that our definition of the error measure can easily be extended to involve further expert knowledge about the local influence of experimental parameters, such as particle seeding density, on the error of the cross-correlation method.

2.3 Joint Optimisation

In section 2.1 and 2.2 we proposed two concepts disregarding their dependencies by assuming that the window shape is fixed during correlation respectively that a vector field is known to estimate the error caused by spatial displacement variations. Now we combine both by defining a bi-level optimisation problem,

$$\min_u C(u, \Sigma) \tag{4}$$

$$\text{with } \Sigma(x) \in \arg\min_{\Sigma \in \mathcal{S}_{++}^2} E(\Sigma, u, x), \ \forall x \in \Omega . \tag{5}$$

The top-level optimisation estimates the displacements u at all positions x in the image by maximising the correlation terms. The window shapes are adapted in the underlying optimisation problems that constrains each $\Sigma(x)$ to a minimum of the error estimation function E and again depends on u.

3 Discretisation and Optimisation

3.1 Data and Variable Discretisation

The discrete input data g_1, g_2 is assumed to be sampled at a regular grid Y with grid spacing a_y and is stored in a cubic spline representation which yields a two times continuously derivable representation. We use an efficient implementation based on [5] to evaluate the function g, its gradient ∇g and its second derivatives $\nabla^2 g$ also at non-integer positions. Grey values of g_1 and g_2 are shifted beforehand so they have a mean of zero each. Values outside the image domain are defined to be zero.

Displacement and window shape variables are discretised on a separate regular grid $X \subset \Omega$ with spacing a_x which is typically chosen to be coarser than Y. We denote the variables located at the coordinates $x_i \in X$ as $u_i := u(x_i)$ respectively $\Sigma_i := \Sigma(x_i)$. For the discretisation of the integral expressions in (2) and (3) we use simple, piecewise constant element functions

$$v_i(x) := \begin{cases} 1 & \text{if } \|x - x_i\|_\infty \leq \frac{1}{2}a \\ 0 & \text{otherwise} \end{cases} ,$$

which incorporate a discrete basis for functions $\mathbb{R}^2 \mapsto \mathbb{R}$ if arranged on a regular grid with spacing a. The discrete versions of the objective functions then read

$$C(u, \Sigma) = a_x^2 \sum_{x_i \in X} a_y^2 \sum_{y_j \in Y} w(y_j - x_i, \Sigma_i) g_1 \left(y_j - \frac{1}{2} u_i \right) g_2 \left(y_j + \frac{1}{2} u_i \right)$$

$$\text{and } E(\Sigma, u, x) = a_x^2 \sum_{x_i \in X} w(x_i - x, \Sigma) e(x, x_i, u) + \frac{\sigma^2}{2\pi\sqrt{\det \Sigma}} \ .$$

The derivative of the first function with respect to u_i simplifies to

$$\nabla_{u_i} C(u, \Sigma) = a_x^2 a_y^2 \sum_{y_j \in Y} w(y_j - x_i) \nabla_{u_i} \left(g_1 \left(y_j - \frac{1}{2} u_i \right) g_2 \left(y_j + \frac{1}{2} u_i \right) \right) \ .$$

For speed up evaluation is restricted to a rectangular area that encloses all positions with $w(x, \Sigma) \geq 10^{-3}$.

3.2 Optimisation

The joint formulation (4)–(5) is a bi-level optimisation problem with non-convex function in both layers. First we remove the explicit constraints $\Sigma_i \succ 0$ of the second level by including them into the objective function as a logarithmic barrier function and define

$$\tilde{E}(\Sigma, u, x) := E(\Sigma, u, x) - \mu \log \det \Sigma \ ,$$

with some constant $\mu > 0$, here chosen as $\mu = 10^{-2}$. In the same manner it is possible to incorporate additional upper and lower bounds on the window size. We relax the optimisation objectives by considering only the first order optimality conditions,

$$\nabla_u C(u, \Sigma) = 0 \tag{6}$$

$$\nabla_{\Sigma_i} \tilde{E}(\Sigma_i, u, x_i) = 0 \ , \quad \forall x_i \in X \ . \tag{7}$$

An initial solution is iteratively improved by updating the variables in parallel.

```
initialise Σ^(0), u^(0), k ← 0
repeat
    k ← k + 1
    for each x_i ∈ X
        u_i^(k+1) ←update(C(·, Σ_i^(k), x_i), u_i^(k), λ_u)
        Σ_i^(k+1) ←update(Ẽ(·, u^(k), x_i), Σ_i^(k), λ_Σ)
until convergence
```

The update step is a Levenberg-Marquardt method that just as the Newton step method involves first and second order derivatives of the objective function $f(x)$, but additionally steers the step length by the parameters $\lambda > 0$.

```
function update(f(x),x₀,λ)
  H ← ∇²f(x₀), g ← ∇f(x₀)
  d ← −(H + λI)⁻¹g
  choose α ∈ [0,1), so that f(x₀ + αd) ≤ f(x₀)
  return x₀ + αd
```

However due to the non-convexity of C and E, a variable assignment that fulfils equations (6)-(7) does not necessarily incorporate a local optima for (4)-(5). For this reason we line search along the step direction to reduce the value of the objective functions and to avoid local maxima and saddle points.

3.3 Multiscale Framework

The optimisation method described in the previous section finds a local optimum which however might be far from a globally optimal solution. Our approach avoids most of them by plugging the single-level optimisation into a multiscale framework. Figure 2 illustrates how this allows to circumnavigate the many suboptimal positions in the correlation function of two noisy images.
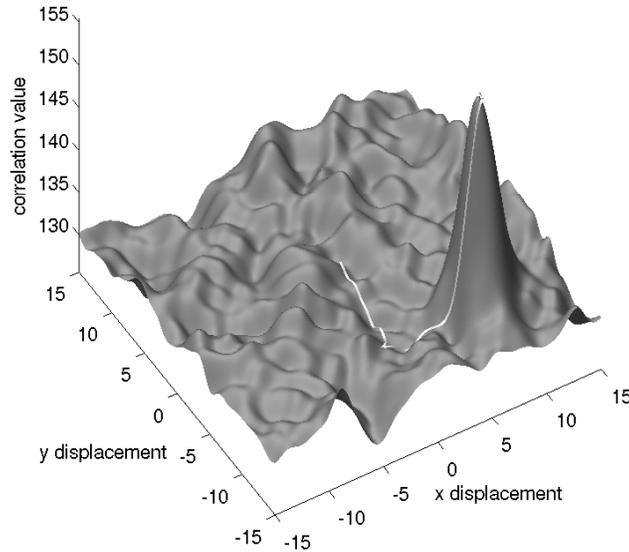

Fig. 2. Optimisation of the non-convex correlation function of noisy image data: Value of $-C(u, \Sigma, x)$ over u and evolution of the displacement over several iterations (bright line, starting in $(0, 0)$); due to the multiscale framework the method does not get stuck in a local optima but finds the correct solution in $(+8, -8)$.

In order to gain a coarse-to-fine representation of the image frames, the data is repetitively low-pass filtered and sub-sampled. When moving in opposite direction within the resolution pyramid, the variables are sampled to a finer grid and

used as initialisation for the next iteration steps. We use cubic spline interpolation both for the re-sampling of data and vector variables. For the window shape variables, bi-linear interpolation implicitly conserves the constraint $\Sigma \in \mathcal{S}_{++}^2$.

4 Experiments and Discussion

4.1 PIV Evaluation Data Set

Our first experiment is based on an evaluation data set designed to verify the ability of correlation-based motion estimation methods to cope with large gradients in the vector field. Case A4 of the PIV Challenge 2005 [6] contains an area named *1D Sinusoids I* which consist of two synthetic particle images each 1000 by 400 pixels in size. Their vertical displacements vary sinus-like with different wavelength (400 pixels on the left down to 20 on the right) while the horizontal component is zero everywhere, see Fig. 3(a)-(b).

Six respectively two scale levels were used for the fixed-window experiments and experiments with window adaptation. The scale factor between two successive levels is $\sqrt{2}$. The parameters in (3) were set to $\sigma = 20$ and $e_{\mathrm{out}} = 10$. Windows were constrained so that their 50%-level curve lies within a radius of about 63 pixels, i.e. $w(x, \Sigma) < 0.5$ for all $\|x\|_2 \geq 63$. The maximum displacement in data is about 2.7 pixels and 1.2 pixels in average.

The results (see Fig. 3(c)) show that pure correlation with fixed window shapes can capture the main structures of the images but fails to accurately estimate the vector fields especially in the presence of large displacements gradients. However when used as initialisation to adaptive correlation we obtained a precise reconstruction (see Fig. 3(d)) of the displacement field, even for the smallest wavelength.

4.2 Real Data

In order to test the ability of our approach to cope with noise and disturbances in real-life turbulent data we applied it to an image pair from a wind tunnel experiment (wake behind a cylinder, see [7]). Figure 4(a) shows the resulting vector field.

The multiscale framework used eight scale levels and a scaling factor of $\sqrt{2}$. Window adaptation parameters were chosen as $\sigma = 10$ and $e_{\mathrm{out}} = 10$. The window radius was constrained to a maximum of about 32 pixels. The average measured displacement is about seven pixels.

Although we used neither regularisation terms nor data validation algorithms (e.g. median filter) on the vector field we obtained a smooth solution without outliers. Figure 4(b)-4(d) demonstrate how the window shapes align themselves along areas of equal displacements and avoid gradients as intended by the design of the window adaptation criterion.
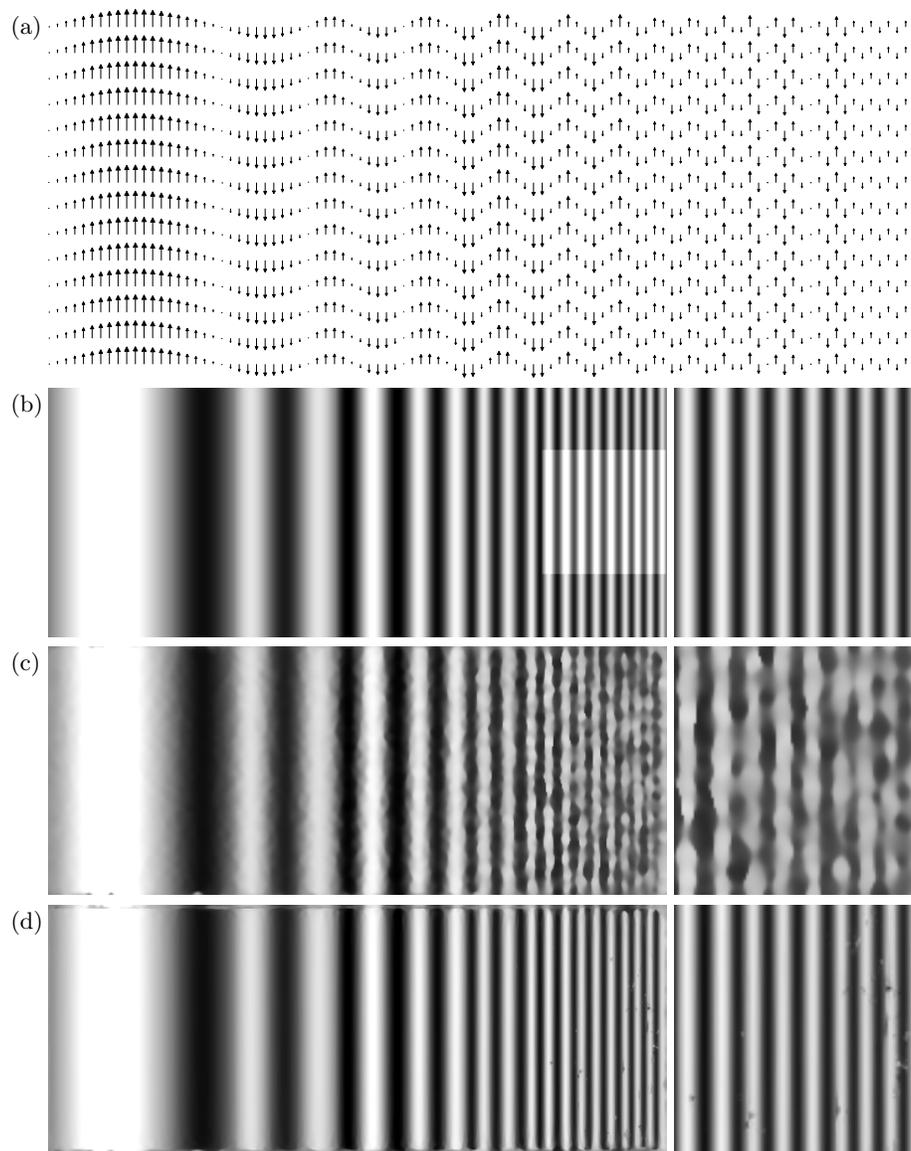
Fig. 3. Experiments with synthetic data: (a) Groundtruth vector field (sub-sampled). (b) Vertical component (mapped to grey-values: bright = up, dark = down) of the groundtruth vector field and highlighted detail, also shown on the right. (c) Correlation with fixed window shape; estimates inevitably deteriorate at small wavelengths. (d) Joint correlation and window adaptation can significantly improve accuracy despite spatially variant wavelengths.
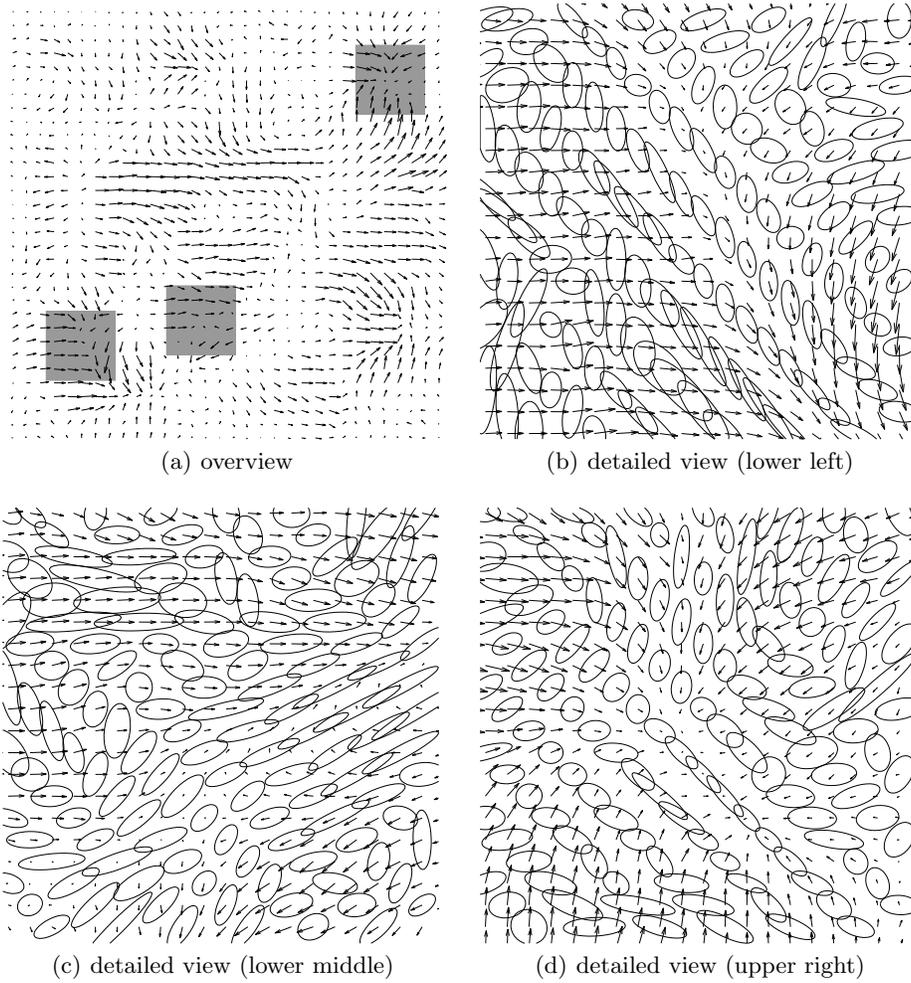
(a) overview   (b) detailed view (lower left)

(c) detailed view (lower middle)   (d) detailed view (upper right)

Fig. 4. Experiments with real data (wake behind a cylinder): (a) Results (sub-sampled) of the correlation approach with window adaptation. The dark highlighting marks the area of the (b)-(d) detailed views of the vector field with some adapted windows (contour line: $w(x, \Sigma) = 0.5$). Note that each window propagates into regions of homogeneous movement and perpendicular to gradients in the vector field and not necessarily along the flow.

5 Conclusion and Further Work

Conclusion. We proposed an approach to fluid flow estimation based on the continuous optimisation of the cross-correlation measure. The expected estimation error for the choice of the correlation window shape is modelled and minimised in order to adapt the windows to displacement gradients. Both subproblems were combined in a bi-level optimisation problem. A multiscale gradient-based approach was described that continuously searches for both optimal displacements and window shapes.

Experiments with synthetic and real data showed that the approach can cope with large displacements and disturbances typical for real fluid flow experiments. The adaptation of window shapes by means of the error expectation model leads to meaningful results in the presence of displacement gradients and improved error significantly in the PIV-Challenge data set.

Further Work. Our approach is the origin for further potential investigations: Due to its variational formulation it is possible to extend the correlation term to estimate also affine displacements within the window and to involve physical priors, such as incompressibility.

Further expert knowledge and statistical information can be incorporated into window adaptation criterion to improve the estimation accuracy. Also more complex shapes for the weighting function should be considered. A comparison to state-of-the-art correlation implementations is planned.

References

1. Raffel, M., Willert, C., Kompenhans, J.: Particle Image Velocimetry. Volume 2. Springer-Verlag, Berlin (2001)
2. Lucas, B., Kanade, T.: An Iterative Image Registration Technique with an Application to Stereo Vision. In: Proceedings of DARPA Imaging Understanding Workshop. (1981) 121–130
3. 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) (2005) 211–231
4. Ruhnau, P., Stahl, A., Schnörr, C.: Variational Estimation of Experimental Fluid Flows with Physics-based Spatio-temporal Regularization. Measurement Science and Technology **18** (2007) 755–763
5. Unser, M., Aldroubi, A., Eden, M.: B-Spline Signal Processing: Part II—Efficient Design and Applications. IEEE Transactions on Signal Processing **41**(2) (1993) 834–848
6. Stanislas, M., Okamoto, K., Kähler, C., Westerweel, J.: Third International PIV-Challenge. (2006, to be published in Exp. in Fluids)
7. Carlier, J., Wieneke, B.: Deliverable 1.2: Report On Production and Diffusion of Fluid Mechanic Images and Data. Activity Report, European Project FLUID Deliverable 1.2, Cemagref, LaVision (2005)