

Variational Optical Flow Estimation for Particle Image Velocimetry

P. Ruhnau¹, T. Kohlberger¹, H. Nobach², C. Schnörr¹

¹Department of Mathematics and Computer Science
Computer Vision, Graphics, and Pattern Recognition Group
University of Mannheim, 68131 Mannheim, Germany
{ruhnau,tiko,schnoerr}@uni-mannheim.de
www.cvgpr.uni-mannheim.de

²Chair of Fluid Mechanics and Aerodynamics
Darmstadt University of Technology
64287 Darmstadt, Germany
holger.nobach@nambis.de
www.sla.maschinenbau.tu-darmstadt.de

Abstract

We introduce a novel class of algorithms for evaluating PIV image pairs. The mathematical basis is a continuous variational formulation for globally estimating the optical flow vector fields over the whole image. This class of approaches has been known in the field of image processing and computer vision for more than two decades but apparently has not been applied to PIV image pairs so far.

We pay particular attention to a multiscale representation of the image data so as to cope with the quite specific signal structure of particle image pairs. The experimental evaluation shows that a prototypical variational approach competes in noisy real-world scenarios with two alternative approaches especially designed for PIV-sequence evaluation. We outline the potential of the variational method for further developments.

1 Introduction

Particle Image Velocimetry (PIV) is an important and active research field concerned with the quantitative investigation of fluids by imaging techniques [15]. The prevailing technique underlying most computational approaches to estimating the motion of fluids from corresponding image pairs is based on the correlation of local interrogation windows in subsequent frames. However, despite the success of this technique and numerous investigations of improvements (which are summarized in [16]), it suffers from some fundamental limitations:

- The partitioning of the image by interrogation areas must not be too fine in order to reliably detect correlation peaks. This unavoidably limits the spatial resolution of the estimated motion vector field.
- The size of the interrogation areas determines a spatial scale at which the variation of motion vector fields is (tacitely) assumed to be negligible. This assumption is too strong and violated in many relevant situations.
- Motion estimation is carried out regardless of spatial context. As a consequence, prior knowledge about spatial flow structures cannot be exploited during estimation, and missing motion estimates in image regions where a correlation analysis yields no reliable estimates have to be heuristically inferred in a post-processing step.

The objective of this paper is to introduce a novel class of motion estimation approaches to the PIV-community. The prototypical approach of this class is due to Horn and Schunck [8], and many corresponding approaches have been analyzed in the fields of image processing and computer vision. For a review and extensions we refer to [21]. In this paper, we focus on the prototypical approach [8] along the analysis presented in [17]. The variational framework and corresponding features will be presented in sec. 2. Next, we carefully examine features of a coarse-to-fine implementation thus taking into account the specific gray value-functions induced by particles in PIV image pairs (sec. 3). Numerical experiments for benchmark image pairs and a comparison with two alternative approaches especially designed for PIV-sequence evaluation will be presented in sec. 4. We conclude in sec. 5 by indicating extensions of the prototypical approach within the variational framework.

2 Variational Approach

2.1 Basic Assumptions and Constraints

Let $f(x, y, t)$ denote the gray value recorded at location $(x, y)^\top$ and time t in the image plane. A basic assumption underlying most approaches to motion estimation is that f is conserved, that is the change of $f(x, y, \cdot)$ at location $(x, y)^\top$ is due to a movement of $f(x, y, t)$ to the location $(x + u\Delta t, y + v\Delta t)^\top$ during a time interval Δt :

$$f(x + u\Delta t, y + v\Delta t, t + \Delta t) = f(x, y, t). \quad (1)$$

Integrating over the entire image domain Ω yields:

$$\int_{\Omega} [f(x + u(x, y)\Delta t, y + v(x, y)\Delta t, t + \Delta t) - f(x, y, t)]^2 dx dy \quad (2)$$

From the viewpoint of variational analysis and algorithm design, formulation (2) is less favourable because the dependency on u and v is highly non-convex. A common way around this difficulty is (i) to further simplify the objective function so as to obtain a mathematically tractable problem, and (ii) to apply the resulting variational approach to a multi-scale representation of the image data f (see sec. 3).

$$\begin{aligned} f(x + u\Delta t, y + v\Delta t, t + \Delta t) \\ \approx f(x, y, t) + \partial_x f(x, y, t)u\Delta t + \partial_y f(x, y, t)v\Delta t + \partial_t f(x, y, t)\Delta t \\ = f(x, y, t) + \nabla f(x, y, t) \cdot \begin{pmatrix} u \\ v \end{pmatrix} \Delta t + \partial_t f(x, y, t)\Delta t \end{aligned} \quad (3)$$

where the spatial and temporal derivatives of f can be estimated locally using FIR filters (see sec. 4.2.). Using (3), the objective function (2) becomes:

$$\int_{\Omega} [\nabla f \cdot \begin{pmatrix} u \\ v \end{pmatrix} + \partial_t f]^2 dx dy \quad (4)$$

Note, that this objective function now depends *quadratically* on the two *functions* $u(x, y)$ and $v(x, y)$ which is much more convenient from the mathematical viewpoint! So far, the transition to a continuous setting has led us to the formulation (4) which has to be minimized with respect to arbitrary functions u and v . Clearly, this problem is not well-posed as yet because *any* vector field with components $\nabla f \cdot (u, v)^\top = -\partial_t f, \forall x, y$, is a minimizer. However, rather than to consider vector fields which are piecewise constant within interrogation areas, we merely rule out too irregular vector fields by additionally minimizing the magnitudes of the spatial gradients of u and v :

$$J(u, v) = \int_{\Omega} \left\{ [\nabla f \cdot \begin{pmatrix} u \\ v \end{pmatrix} + \partial_t f]^2 + \lambda(|\nabla u|^2 + |\nabla v|^2) \right\} dx dy, \quad 0 < \lambda \in \mathbb{R}. \quad (5)$$

In this paper, we regard λ as a user-parameter. In view of the limitations mentioned in sec. 1, we point out the following features of the variational approach (5):

- The approach is formulated in terms of *functions* u and v and hence, by definition, provides motion estimates $(u(x, y), v(x, y))^\top$ at *any* point $(x, y)^\top \in \Omega \subset \mathbb{R}^2$.
- Spatial variation of u, v is merely constrained by a *global* penalty term (i.e. the second term in (5)). Accordingly, the motion field $(u, v)^\top$ may exhibit spatial variations of different strengths depending on the evidence provided by the spatio-temporal image sequence data f .
- The approach is intrinsically non-local and allows to incorporate spatial context in a mathematically convenient way by means of functionals depending on u, v and corresponding derivatives.

2.2 Optimization Problem and Discretization

Under mild conditions with respect to the image sequence data f it can be shown [17] that the functional (5) is strictly convex. As a consequence, the *Finite Element Method (FEM)* can fully be exploited, as will be sketched next. For details we refer to [6, 19]. The unique globally minimizing vector field $(u(x, y), v(x, y))^T$ is determined by the variational eqn.

$$a((u, v)^T, (\tilde{u}, \tilde{v})^T) = b((\tilde{u}, \tilde{v})^T), \quad \forall \tilde{u}, \tilde{v}, \quad (6)$$

where

$$a((u, v)^T, (\tilde{u}, \tilde{v})^T) = \int_{\Omega} \left\{ \begin{pmatrix} u \\ v \end{pmatrix} \cdot \nabla f \nabla f \cdot \begin{pmatrix} \tilde{u} \\ \tilde{v} \end{pmatrix} + \lambda (\nabla u \cdot \nabla \tilde{u} + \nabla v \cdot \nabla \tilde{v}) \right\} dx dy, \quad (7)$$

$$b((\tilde{u}, \tilde{v})^T) = - \int_{\Omega} \partial_t f \nabla f \cdot \begin{pmatrix} \tilde{u} \\ \tilde{v} \end{pmatrix} dx dy. \quad (8)$$

The simplest discretization is obtained by choosing a regular triangulation of the image domain Ω and attaching to each pixel position a piecewise linear basis function $\phi(x, y)$ for each function $u, v, \tilde{u}, \tilde{v}$. Indexing each pixel position (k, l) by $1, 2, \dots, N$ we thus have

$$u(x, y) = \sum_{i=1}^N u_i \phi_i(x, y),$$

and similarly for v, \tilde{u}, \tilde{v} . This leads to a linear system of size $2N \times 2N$ which is sparse and positive definite. Thus u, v can be conveniently computed by some corresponding iterative solver [7].

3 Coarse-to-Fine Motion Estimation

The accuracy of motion estimation critically depends on the magnitude of image motion. In fact, depending on the spatial image frequency, very large motions even may cause aliasing along the time frequency axis. Due to the Nyquist-condition $|\omega_t| < \pi$ (with $\omega_t := \omega_x u$), only motions up to $|u| < \pi/\omega_x$ are correctly represented by samples of the signal.¹ Faster motions lead to aliasing. In other words, for a fixed global velocity, spatial frequencies moving more than half of their period per frame cause temporal aliasing. In practice, this upper bound has to be lowered because derivatives of the signal can only be robustly estimated in connection with low-pass filtering.

As a remedy, we first compute a coarse motion field by using only low spatial frequency components and “undo” the motion, thus roughly stabilizing the position of the image over time. Then the higher frequency subbands are used to estimate optical flow on the warped sequence. Combining this “optical flow correction” with the previously computed optical flow yields a refined overall optical flow estimate. This process may be repeated at finer and finer spatial scales until the original image resolution is reached [11, 20]. A standard technique for generating multi-scale representations in this context is to construct an image pyramid (Fig. 1) by recursively applying lowpass filtering and subsampling operations. Note that the images at different scales are represented by different sampling rates. Thus, the same derivative filters may be used at each scale and we do not have to design multiple derivative filters, one for each different scale. Let us define the pyramid representation of a generic image f of size $n_x \times n_y$. Let $f^0 = f$ be the “zeroth” level image. This image is essentially the highest resolution image (the raw image). The image width and height at that level are defined as $n_x^0 = n_x$ and $n_y^0 = n_y$. The pyramid representation is then built in a recursive fashion: Compute f^1 from f^0 , then compute f^2 from f^1 , and so on Let $k = 0, 1, 2, \dots, L-1$ be a generic pyramidal level, and let f^k be the image at level k . n_x^k and n_y^k denote the width and the height of f^k . First the lowpass filter $[1/4 \ 1/2 \ 1/4] \times [1/4 \ 1/2 \ 1/4]^T$ is used for image anti-aliasing before image subsampling. Then a bilinear interpolation performs the adaption to the new coarser grid, as every new vertex is located exactly in the middle of four finer vertices (if the respective image size is even-numbered, cmp. Fig. 2). This procedure results in a convolution mask of





Figure 1: Image Pyramid: Each level in the pyramid is a subsampled version of the level below convolved with a Gaussian filter.

Figure 2: Image Pyramid: Location of the vertices in the respective levels.

$[1/8 \ 3/8 \ 3/8 \ 1/8] \times [1/8 \ 3/8 \ 3/8 \ 1/8]^T$. In the first step the optical flow between the top level images f_1^{L-1} and f_2^{L-1} (lowest frequency images) is computed, using the variational approach of sec. 2. The computed coarse-level flow field must then be projected onto the next finer pyramid level. This flow field estimate is used to warp the second image towards the first image:

$$\mathcal{W}\{f_2^{L-1}, d^{L-1}\}(x, y, t + \Delta t) = f_2^{L-1}(x - u\Delta t, y - v\Delta t, t + \Delta t), \quad d^{L-1} = \begin{pmatrix} u \\ v \end{pmatrix} \quad (9)$$

At pyramid level $L - 2$, we compute a new and finer flow field between the images f_1^{L-2} and $\mathcal{W}\{f_2^{L-2}, d^{L-1}\}$. While the expression to be minimized is analogous to (5), the first-order Taylor series expansion is performed around $(x + d^{L-1}(x), t + 1)$. This results in the cost functional:

$$J(u, v) = \int_{\Omega} \left\{ [\nabla f \cdot \begin{pmatrix} u \\ v \end{pmatrix} + \partial_t f]^2 + \lambda (|\nabla(u + d_u^{L-1})|^2 + |\nabla(v + d_v^{L-1})|^2) \right\} dx dy, \quad 0 < \lambda \in \mathbb{R}. \quad (10)$$

The unique flow field minimizing (10) is the correction-field Δd^{L-2} of the coarser flow d^{L-1} . To obtain the overall flow field d^{L-2} at level $L - 2$ we have to add the coarse motion d^{L-1} and the correction field Δd^{L-2} :

$$d^{L-2}(x, y) = d^{L-1}(x, y) + \Delta d^{L-2}(x, y) \quad (11)$$

This correction process is repeated for each level of the pyramid until the finest pyramid level d^0 has been reached. In the experimental evaluation sec. below, we will refer to this approach as Horn & Schunck Multi-Resolution (**H&S R**). So far, we have introduced a *dyadic* pyramid structure which is equivalent to using lowpass filters with bandwidths $\frac{\Omega}{2^{L-1}}, \frac{\Omega}{2^{L-2}}, \dots, \frac{\Omega}{2^1}, \frac{\Omega}{2^0}$ combined with subsampling. Now we introduce additional filters that slice the bandwidth into even smaller pieces, e.g. $\Omega/4, 3/8\Omega, \Omega/2, 3/4\Omega, \Omega$. In order to implement these extra steps which do not fit into the dyadic pyramid structure, we apply at each pyramid level pre-filters when estimating derivatives: The lower the cut-off frequency of the pre-filter, the more the particles seem to melt down and form a smooth gray value structure. A coarse motion estimate can reliably be computed using this structure. Then, we update and refine the motion field (in the same way as described in detail for the multi-resolution case) using the less low-pass filtered image derivatives. For the experiments in this paper, we use nine scale-space levels and thus nine different filters with cut-off frequencies of $\frac{\pi}{2}, \frac{9}{16}\pi, \frac{5}{8}\pi, \frac{11}{16}\pi, \frac{3}{4}\pi, \frac{13}{16}\pi, \frac{7}{8}\pi, \frac{15}{16}\pi, \pi$. An inverse Fourier Transform yields the filter coefficients. Low pass filtering with cut-off frequencies below $\pi/2$ is not necessary since this is what the anti-aliasing filter of the preceding lower resolution level has already done. Below, we will refer to this combined approach as Horn&Schunck Multi-Resolution + Multi-Scale (**H&S R+S**).

4 Experimental Evaluation

In this sec., we report comparisons of the variational approach with two other approaches for various data sets. Before discussing the results in sec. 4.4 below, we first describe the data sets used for the

¹Without loss of generality we assume sampling rates $\Delta x = \Delta t = 1$.




Figure 3: Quénnot image pair: synth. particle image




Figure 4: Exact velocity field (av. 7.58 px./frame)

comparison (sec. 4.1), the alternative approaches (besides the variational approach) and corresponding parameter setting (sec. 4.2) and quantitative error measures (sec. 4.3).

4.1 Data

The experimental evaluation was carried out on the basis of the following data sets:

- **Synthetic Data:** The “Quénnot image pair” was introduced in [14] and is available on the internet. The analyzed velocity field (av. velocity = 7.58 px./frame) is taken from a numerical solution obtained for two-dimensional flow around a pair of cylinders (Fig. 3). We examined different test cases:
 - **Perfect:** “Perfect” case means that the second image was computer-generated from the first image and the target flow field.
 - **Mixed N%:** The specified percentage of noise was superimposed. Additionally, the specified percentage of particles was randomly removed and the same amount of particles was randomly added.




Figure 5: Real world image




Figure 6: Estimated flow field (variational approach)

- **Real world image:** Fig. 5 shows a corresponding image from a time-resolved PIV measurement of periodically vortices in the transitional cylinder wake [3, 23]. The mean displacement is about 9 px./frame and the maximum displacement about 16 px./frame.

4.2 Approaches and Parameter Settings

The data sets described above were evaluated using the following approaches and parameter settings:

- **Variational approach:** The spatial (∇f) and temporal ($\partial_t f$) derivatives were estimated using derivative of Gaussian filters of size five at every point in the image domain. In a first series of experiments (**H&S R**) five resolution levels were used, in a second series of experiments (**H&S R+S**) a setup of five resolution levels and nine scale space levels on every resolution level was chosen. For the Quénnot image pair computations, the smoothness parameter λ was set to $7 \cdot 10^{-4}$ in the H&S R case and to $7 \cdot 10^{-3}$ in the H&S R+S case. For the real world image pair only H&S

R+S computations were performed. The smoothness parameter λ was also set to $7 * 10^{-3}$ in these rows of experiments. The gray values were scaled in each case to the interval $[0, 1]$.

- **DIPV approach:** For comparison we took the error measures of the classical 2D FFT based digital particle image velocimetry (DPIV) method from [14] in the synthetic test cases. Two different interrogation window sizes were applied: 32×32 pixel (**DPIV 32**) and 48×48 pixel (**DPIV 48**). We analyzed the “cylinder wake” real world image pair using a hierarchical DPIV approach with an interrogation window size beginning with 512×512 pixels and ending up with 64×64 pixels with window-shifting and peak-height validation.
- **ODP2 approach [13]:** We considered also the results of a dynamic programming based optical flow technique. This approach transforms the two-dimensional correspondence problem to a sequence of one-dimensional search problems. It has been successfully applied to PIV in [14, 12].

4.3 Error Measures

As quantitative error measures we computed the angular error (between correct and computed motion vectors) as defined in [1] along with its standard deviation as well as the mean velocity error (L_1 norm of the difference between the correct and computed velocities in px./frame). The error measure was computed for the whole image except for the inner circular regions corresponding to the cylinders.

4.4 Numerical Results and Discussion

4.4.1 Quénot Image Pair

Table 1 summarizes the error measures and their standard deviation (\pm)². Furthermore, typical execution times of the respective algorithms are indicated. Note that DPIV yields a sparse vector field whereas

		DPIV32	DPIV48	ODP2	H&S R	H&S R+S
Perfect	angle	5.95 ± 13.9	9.35 ± 18.3	1.23 ± 2.24	2.32 ± 3.69	1.42 ± 2.16
	disp	0.55 ± 0.94	0.87 ± 1.46	0.13 ± 0.10	0.39 ± 0.66	0.19 ± 0.18
Mixed 5%	angle	6.40 ± 14.4	9.59 ± 19.0	1.77 ± 2.87	3.04 ± 4.38	1.54 ± 2.34
	disp	0.60 ± 1.12	0.86 ± 1.51	0.20 ± 0.13	0.52 ± 0.80	0.21 ± 0.30
Mixed 10%	angle	10.2 ± 19.6	11.3 ± 20.8	4.30 ± 11.7	4.39 ± 5.89	1.81 ± 2.74
	disp	0.91 ± 1.89	0.93 ± 1.66	0.57 ± 1.71	0.78 ± 1.07	0.26 ± 0.46
Mixed 20%	angle	40.8 ± 34.5	38.3 ± 29.7	6.15 ± 9.01	9.33 ± 9.93	2.96 ± 4.23
	disp	3.73 ± 4.39	2.49 ± 3.19	0.74 ± 0.52	2.03 ± 2.13	0.39 ± 0.53
Time		10 min	10 min	20 min	16 sec / 2 sec	2 min / 15 sec

Table 1: Angular error and absolute displacement error. Best performance for every setting is marked in bold.

both ODP2 and H&S compute dense vector fields. All of the tested algorithms are (in varying degrees) sensitive to superimposed noise. In the case of DPIV, extending the interrogation window size increases the robustness to noise while decreasing the accuracy at the same time. However, irrespective of the window size the performance of DPIV is much worse than the performance of the other approaches.

We realize that H&S R+S provides much better results than H&S R in all test cases. This had to be expected because temporal aliasing as well as linearization errors due to eqn. (3) are suppressed by additional scale space computations. Fig. 7 shows the results for the “Mixed 20%” case. One can see that the highest estimation errors are reached at the borders of the cylinders. The smoothness term penalizes the discontinuities at these locations and smoothes over the discontinuities. The error at regions close to the left cylinder is the highest because of the high velocity of the fluid. In sec. 5 we will point out possible solutions to this problem such as the insertion of border conditions and higher order regularization.

ODP2 provides the best result for the “perfect” test case. However, it is much less robust to noise than H&S: While the error measures of the variational coarse-to-fine approach are slightly higher in the perfect case (cf. table 1), this changes with the presence of noise. The error for the ODP2 approach then rises much faster so that, for noise rates of 5 % and above, the H&S R+S approach provides better results.

²Error measures for the three algorithms not implemented by the authors were taken from [14], the execution times from [5].




Figure 7: Results for the Quénot image pair “Mixed 20%”. Estimated flow field with H&S R+S (left), absolute displacement error (right).

This gap rises even more for higher amounts of noise. Since noise is always present in real world images pairs, we expect the H&S algorithm to perform better than both the DPIV and the ODP2 approaches. When we use a preconditioned conjugate gradient method to solve the H&S system matrices, the execution time of our algorithm is about 16 sec for H&S R and 2 min for H&S R+S (when choosing a residual error of 10^{-4} as a stopping criterion). Using a multi-grid approach [2, 7] to solve the linear systems, the computation time of H&S R is approx. 2 sec, while the H&S R+S computation takes about 15 sec on an up-to-date computer. Information about the different multi-grid cycles and stopping criterions can be taken from [4]. Even real time operation can be achieved through parallelization using domain decomposition [10].

4.4.2 Results on a Real World Image Pair

Figures 8 and 9 show the results for the real world image pair (“cylinder wake”) computed with H&S R+S and DPIV, respectively. One can clearly see that the variational approach resembles the true motion




Figure 8: Dense vector field computed with the variational approach.

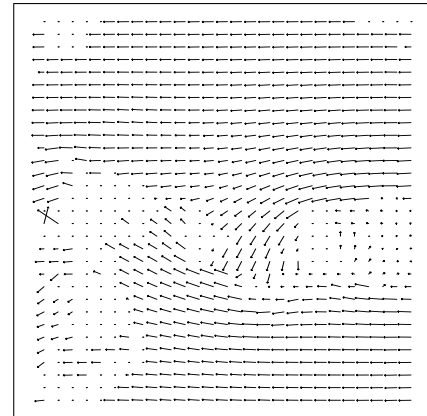


Figure 9: None-dense vector field computed with DPIV.

field much better than the cross-correlation approach. At regions with abruptly changing motion (i.e. the turbulence emerging behind the cylinder in the middle of the image), the DPIV method is not able to accurately determine the velocity field. This is mainly due to the limited spatial resolution which leads to a violation of the assumption of a constant velocity inside interrogation windows at these locations. The statistical character of correlation-based processing, however, prohibits the use of smaller interrogation windows. Furthermore, in regions dominated by out-of-plane velocities (i.e. at the left border of the image), the cross correlation approach fails as well: Since no global velocity information is used, the probability of outliers is markedly increased at these locations, hence a valid flow field cannot be computed.

5 Conclusions and Further Work

We have successfully applied a prototypical variational optical flow estimation approach to Particle Image Velocimetry. The novel approach outperforms the standard cross correlation methods and computes dense motion fields. We also compared our approach to ODP2, a sophisticated optical flow technique that is often used for PIV code validation. While ODP2 led to slightly better results for noise-free image

pairs, our approach produces better results for noise levels of 5% and above. As a result, we expect that our variational optical flow estimation approach will perform better in real world routine computations.

A decisive advantage of the variational approach (5) is its potential for further developments. Note that the data term (first term in (5)) was originally designed for images from everyday scenes (e.g. traffic scenes) whose gray value functions f are quite different and more well-behaved than PIV-sequences from the signal processing point-of-view. While we tried to cope with this difficulty by carefully designing a coarse-to-fine representation of the image data (sec. 3), alternative data terms are conceivable, in particular in the critical case of low particle densities. Furthermore, the data term could be expanded to additionally estimate parameters of an illumination model to cope with the fact that local intensity changes are common in PIV image sequences. Similarly, various extensions of the simple smoothness term in (5) are possible, like spatio-temporal regularization [22], div-curl-shear regularization [18] or non-quadratic discontinuity-preserving regularization [21], for instance. Furthermore, one could add the possibility to specify border conditions for regions where a liquid rinses a solid. Finally, the mathematical formulation leads to sound parallel implementations using off-the-shelf hardware [4, 9].

References

- [1] BARRON, J., FLEET, D., AND BEAUCHEMIN, S., Performance of optical flow techniques. *IJCV*, 39:43–77, 1994.
- [2] BRANDT, A., Multi-level adaptive solutions to boundary-value problems. *Mathematics of Computation*, 31(138):333–390, 1977.
- [3] BREDE, M., LEDER, A., AND WESTERGAARD, C.H., Time-resolved PIV investigation of the separated shear layer in the transitional cylinder wake. In *Proc. 5th Int. Symp. on PIV, (PIV'03), Busan, Korea*, 2003.
- [4] BRUHN, A., WEICKERT, J., FEDDERN, C., KOHLBERGER, T., AND SCHNÖRR, C., Real-time optic flow computation with variational methods. In N. Petkov and M.A. Westenberg, editors, *Proc. Computer Analysis of Images and Patterns (CAIP'03)*, volume 2756 of *LNCS*, 222–229. Springer, 2003.
- [5] CHETVERIKOV, D., Applying Feature Tracking to PIV *IJPRAI*, 17 (4):477–504, 2003.
- [6] CIARLET, P.G., *The Finite Element Method for Elliptic Probl.*. North-Holland Publ. Comp., Amsterdam, 1978.
- [7] HACKBUSCH, W., *Iterative Solution of Large Sparse Systems of Equations*. Springer-Verlag, 1993.
- [8] HORN, B.K.P., AND SCHUNCK, B.G., Determining optical flow. *Artif. Intell.*, 17: 185–203, 1981.
- [9] KOHLBERGER, T., SCHNÖRR, C., BRUHN, A., AND WEICKERT, J., Domain decomposition for parallel variational optical flow computation. In B. Michaelis and G. Krell, editors, *Pattern Recognition, Proc. 25th DAGM Symposium*, volume 2781 of *LNCS*, 196–203. Springer, 2003.
- [10] KOHLBERGER, T., SCHNÖRR, C., BRUHN, A., AND WEICKERT, J., Parallel Variational Motion Estimation by Domain Decomposition and Cluster Computing. In T. Pajdla and J. Matas, editors, *Proc. ECCV 2004*, volume 3024, 205–216. Springer, 2004.
- [11] KRISHNAMURTHY, R., MOULIN, P., AND WOODS, J. Multiscale motion models for scalable video coding, 1996.
- [12] QUÉNOT, G. M., Performance Evaluation of an Optical Flow Technique applied to PIV using the VSJ Standard Images. *Third International Workshop on PIV* 579–584, 1999.
- [13] QUÉNOT, G. M., The orthogonal algorithm for optical flow detection using dynamic programming *Intl. Conf. on Acoustics, Speech and Signal Proc.* 249–252, 1992.
- [14] QUÉNOT, G. M., PAKLEZA, J., AND KOWALEWSKI, T. A., PIV with optical flow. *Exp.Fluids*, 25: 177–189, 1998.
- [15] RAFFEL, M., WILLERT, C., AND KOMPENHANS, J., *PIV*. Springer, 2nd edition, 2001.
- [16] SCARANO, F., Iterative image deformation methods in PIV. *MST*, 13: R1–R19, 2002.
- [17] SCHNÖRR, C., Determining optical flow for irregular domains by minimizing quadratic functionals of a certain class. *Int. J. of Comp. Vision*, 6(1): 25–38, 1991.
- [18] SCHNÖRR, C., Segmentation of visual motion by minimizing convex non-quadratic functionals. In *12th Int. Conf. on Pattern Recognition*, Jerusalem, Israel, Oct 9–13 1994.
- [19] SCHNÖRR, C., Variational methods for adaptive image smoothing and segmentation. In B. Jähne, H. Haußecker, and P. Geißler, editors, *Handbook on Computer Vision and Applications: Signal Processing and Pattern Recognition*, volume 2, 451–484, San Diego, 1999. Academic Press.
- [20] SIMONCELLI, E. P., *Distributed Representation and Analysis of Visual Motion*. PhD thesis, MIT, 1993.
- [21] WEICKERT, J. AND SCHNÖRR, C., A theoretical framework for convex regularizers in pde-based computation of image motion. *IJCV*, 45(3): 245–264, 2001.
- [22] WEICKERT, J., AND SCHNÖRR, C., Variational optic flow computation with a spatio-temporal smoothness constraint. *J. Math. Imaging and Vision*, 14(3): 245–255, 2001.
- [23] WESTERGAARD, C.H., BREDE, M., AND LEDER, A., Time-space analysis of time resolved PIV data. In *Proc. 5th Int. Symp. on PIV (PIV'03), Busan, Korea*, 2003.