

Variational Correlation Approach to Flow Measurement with Window Adaption*

Florian Becker¹, Bernhard Wieneke², Jing Yuan¹, Christoph Schnörr¹

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@lavisoin.de

Abstract Cross-correlation constitutes the state-of-the-art method for measuring motion in PIV data. We present a novel variational approach to correlation-based fluid flow estimation, which determines the displacements via continuous optimisation. A Gaussian shaped window function is used, whose size, eccentricity and orientation can freely be chosen. We adapt the shape to the noise level and displacement gradients by minimising an error function that directly formulates the aim to minimise the motion estimation error. Displacement measurement and window adaption are performed simultaneously in a joint bi-level optimisation problem.

The approach showed its ability to significantly improve accuracy also in the presence of high velocity gradients in synthetic data. Moreover, we tested our method with real PIV images. In comparison to a state-of-the-art cross-correlation implementation, the window adaption leads to smoother displacement fields in regions of homogenous motion, while gradients such as in vertexes are not affected in intensity and general shape.

1. Introduction

Cross-correlation constitutes the state-of-the-art method for measuring displacements in particle image velocimetry (PIV) data. In the recent years, much effort has been put into improving the spatial resolution by replacing the square correlation windows by appropriate alternatives. Nogueira et al. (2005) consider a class of cone-shaped weighting functions and optimise the parameters by means of the amplitude-frequency-response. In Di Florio et al. (2002) a Gaussian shaped window is stretched and rotated along the measured mean displacement. Theunissen et al. (2007) use square windows and locally adapt their size to the signal quality (seeding density) and spatial fluctuations in the flow.

In this work, we present a novel variational approach to correlation-based estimation of fluid-flows. We use a parameterisable Gaussian weighting function for masking the image information. The window orientation, eccentricity and size are adapted by minimising an error function that estimates the influence of displacement gradients, image distortions and window shape on the velocity measurement accuracy. A bi-level optimisation problem is defined that simultaneously adapts the window parameters and estimates the motion field.

Section 2 starts with the description of the proposed variational correlation approach to displacement estimation and the employed weighting function. Afterwards we define the window adaption criterion and comment on its properties. Finally, we combine both the approach to displacement estimation and window adaption into a joint optimisation problem and briefly detail on optimisation and discretisation. In section 3, we verify the algorithm design with synthetic and real PIV data. A conclusions and an overview over further work is given in section 4.

* This work was partially financed by the EC project FLUID (FP6-513663).

2. Approach

2.1. Correlation as Variational Problem

For the continuous definition of our approach we assume that the input data consists of two grey-valued images g_1 and g_2 , which are defined on the coordinates $\Omega \subseteq \mathbb{R}^2$, for example the camera chip area. Outside Ω we define the image function to be zero.

Then we can express the estimation of a displacement $v \in \mathbb{R}^2$ between the images g_1 and g_2 at position $x_0 \in \Omega$ as the maximisation of the correlation measure,

$$c(v, x_0) := \int_{\mathbb{R}^2} w(y - x_0) g_1\left(y - \frac{1}{2}v\right) g_2\left(y + \frac{1}{2}v\right) dy. \quad (1.1)$$

The central differencing interrogation scheme, which shifts the image data by half the displacement forward respectively backward, is second-order accurate in time (see Werely and Meinhardt 2001). Note that in general the whole image data is involved in the measurement due to the integration over all \mathbb{R}^2 . However, we restrict the correlation to a neighbourhood of x_0 using a general weighting function $w(x)$. A common choice is a square window with side length a , i.e.

$$w(x, a) := \begin{cases} 1 & \text{if } -\frac{1}{2}a \leq x_1 \leq +\frac{1}{2}a \text{ and } -\frac{1}{2}a \leq x_2 \leq +\frac{1}{2}a \\ 0 & \text{otherwise} \end{cases}, \quad (1.2)$$

where x_1 and x_2 are the two components of vector x , see Fig. 1(a). In this work, however we define the window function as a non-normalised Gaussian function (see Fig. 1(b)),

$$w(x, \Sigma) := \exp\left(-\frac{1}{2}x^T \Sigma^{-1}x\right),$$

whose shape is parameterised by a symmetric positive definite 2-by-2-matrix Σ . Fig. 2 depicts some possible shapes and demonstrates the three degrees of freedom of the window function: the ratio between length and width of the Gaussian function - denoted as *anisotropy* or *eccentricity* -, *orientation* and the overall *size*. This formulisation allows to *continuously* adapt the correlation support to local requirements with only few parameters.

In contrast to (1.2), this window is not sharply bounded, but has a maximum value of one at its centre of mass ($x = 0$) and steadily tends to zero in all directions. Although in theory $w(x, \Sigma)$ does not vanish anywhere and thus all image data is involved in correlation, in practice it is a very good approximation to consider only those regions where the weight is larger or equal to a boundary, e.g. 0.001.

The support of the window function is $\int_{\mathbb{R}^2} w(x, \Sigma) dx = 2\pi\sqrt{\det \Sigma}$. For illustration, we also give a window with equal support and sharp boundary, which is implicitly defined by the level curve $w(x, \Sigma) = \exp(-1)$ in Fig. 1(c). Whenever we refer to dimensions (e.g. radius) in context with Gaussian windows, we implicitly measure them at this curve.

For now, we assume that we have given a fixed window shape $\Sigma(x)$ everywhere on Ω and define an optimisation problem that simultaneously estimates displacements $u(x)$ at *all* position $x \in \Omega$:

$$\max_{u(x)} C(u(x), \Sigma(x)) \quad \text{with } C(u(x), \Sigma(x)) := \int_{\Omega} c(u(x), \Sigma(x), x) dx \quad (1.3)$$

where $c(v, \Sigma, x)$ is the correlation measure (1.1) extended by the window shape parameter. Note that this formulation allows adding spatial regularisation terms to the objective function and thus incorporate prior knowledge on the vector field, e.g. incompressibility of the flow, see Ruhnau et al (2007).

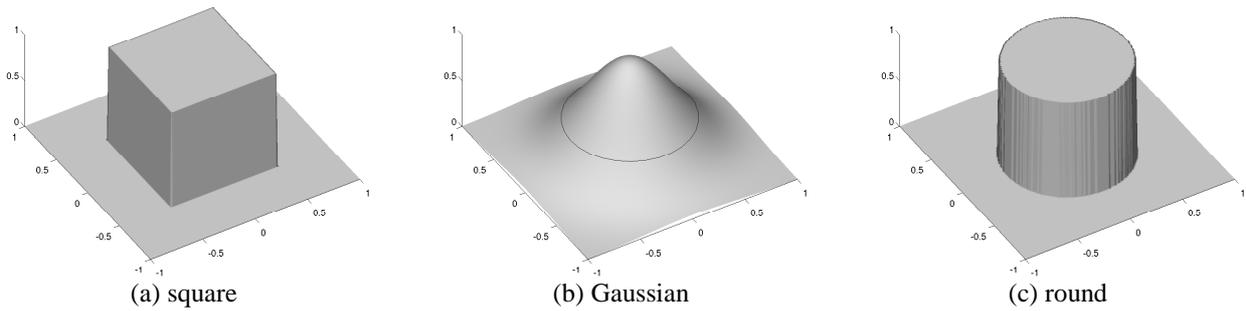


Fig. 1: Three weighting functions $w(x, p)$ with equal support $\int_{\mathbb{R}^2} w(x, p) dx = 1$: (a) square window, (b) Gaussian window with infinite support; the level curve $w(x, \Sigma) = \exp(-1)$ defines the boundaries of a (c) round window with radius $\pi^{-1/2} \approx 0.564$.



Fig. 2: Parameterisation of the Gaussian weighting function: Three degrees freedom and some possible shapes.

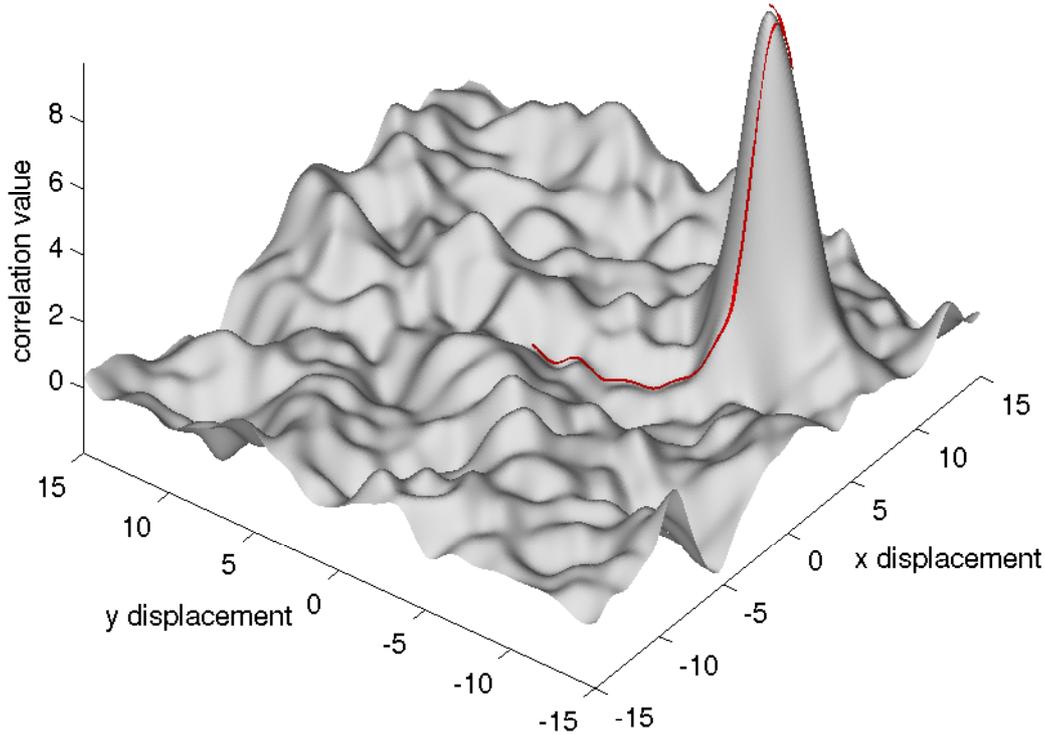


Fig. 3: Surface plot of the correlation function between two noisy PIV images with several local maxima, and the trajectory (red line starting in (0,0)) of the continuous optimisation of the displacement u . Due to the multiscale framework the method does not get stuck in a local optima but finds the correct correlation peak in (+8,-8).

2.2. Window Adaption

In (1.3) we search for the optimal displacements by maximising the cross-correlation function. We implicitly assume that the movement within the considered image patch is uniform, which however does not hold true in most data. The estimated displacement is rather an average of the motion within the correlation window. Thus, large gradients in the vector fields are smoothed out, causing a loss in details. This effect can be avoided by reducing the window size, but with the harm of a smaller area of support and number of particles and thus higher influence of noise.

Hence, for an accurate solution it is important to balance the influence of error sources by adapting the window size and shape. For this purpose, we define a function that models the expected error in the motion estimation process caused by the previously described opposing effects subject to the window parameter:

$$E(\Sigma, u(x), x_0) := \underbrace{\int_{\mathbb{R}^2} w(y - x_0, \Sigma) e(x_0, y, u(x)) dy}_{(a)} + \underbrace{\frac{\sigma^2}{2\pi\sqrt{\det \Sigma}}}_{(b)}$$

$$\text{with } e(x_0, y, u(x)) := \begin{cases} \|u(y) - u(x_0)\|_2^2 & \text{if } y \in \Omega \\ e_{\text{outside}}^2 & \text{otherwise} \end{cases}$$

We presume that a vector field $u(x)$ defined on Ω is known beforehand. Then the first term (a) describes the expected error caused by motion inhomogeneities around the position x_0 . The function $e(x_0, y, u(x))$ measures the deviation from the assumption that the motion is equal to the displacement $u(x_0)$ at position y . For coordinates where no image information is available, i.e. $y \notin \Omega$, we assume a constant error e_{outside}^2 . The resulting value of the first term is the weighted accumulated error, using the same window function $w(x, \Sigma)$ as for the correlation.

The denominator of the second term (b) is exactly the support of the weighting function, i.e. $\int_{\mathbb{R}^2} w(x, \Sigma) dx = 2\pi\sqrt{\det \Sigma}$. Thus, the value of the fraction grows with smaller window size and describes the effect of insufficient large support in displacement estimation. The parameter σ represents the influence of seeding density, out-of-plane-motion (unmatched particles), image sensor noise and other experimental factors on the displacement estimation error. Furthermore we can give a sound derivation for this term: We assume that a displacement measurement of \hat{v} , based on information within the window described by $w(x, \Sigma)$, is a weighted least-square estimation over independent measurements $\bar{v}(x)$:

$$\hat{v} := \arg \min_{v \in \mathbb{R}^2} \int_{\mathbb{R}^2} w(x, \Sigma) \|v - \bar{v}(x)\|_2^2 dx \quad (1.4)$$

We assume that each $\bar{v}(x)$ is an estimation of the true unit displacement v^* , but disturbed by Gaussian noise, i.e. $\bar{v}(x) \sim \mathcal{N}\{v^*, \sigma^2 I\}$, with the standard deviation σ and the identity matrix I . Then it is possible to show that the expected square deviation of \hat{v} from v^* is given by:

$$\mathcal{E}\left\{\|\hat{v} - v^*\|_2^2\right\} = \frac{\sigma^2}{2\pi\sqrt{\det \Sigma}}$$

Note that the second term only describes the error caused by image data disturbances in the presence of *homogenous* motion. The influence of gradients and other inhomogeneities in the vector field is covered by the term (a) only.

The defined energy function can easily be extended to involve further expert knowledge on the influence of experimental parameters on the estimation accuracy. In particular, here we choose a constant value for σ , but it is possible to adjust the parameter locally in order to adapt to varying influences such as inhomogeneous seeding density or image noise level.

The window used for correlation at position x_0 is then chosen from the set of shapes that minimise the defined error function:

$$\Sigma(x_0) \in \arg \min_{\Sigma} E(\Sigma, u(x), x_0)$$

During the optimisation process, the second term of the error function tries to extend the window in all directions with the aim to increase support. In contrast, the first term avoids that the window grows into regions of large displacement deviations but allows covering of areas with motion similar to $u(x_0)$. The higher the parameter e_{outside} is chosen the more the adaption process will avoid areas outside the image definition range.

2.3. Joint Displacement Estimation and Window Adaption

In section 2.1 we presented a variational description of a global displacement estimator with locally parameterisable window shapes. However, we assumed that the parameters $\Sigma(x)$ are known a priori. Subsequently we defined a window adaption scheme that in turn incorporates the knowledge on a displacements field $u(x)$. In order to resolve the interdependency between these two approaches, we formulate the following bi-level optimisation problem:

$$\begin{aligned} & \max_{u(x)} C(u(x), \Sigma(x)) \\ & \text{with } \Sigma(x_0) \in \arg \min_{\Sigma} E(\Sigma, u(x), x_0) \quad \text{at all positions } x_0 \in \Omega \end{aligned} \quad (1.5)$$

A solution of this problem consists of a global displacement estimation $u(x)$ that maximises the correlation measures. The window shapes $\Sigma(x)$ are adapted simultaneously to minimise the error function at the according position and for the determined displacements.

2.4. Optimisation and Implementation

The continuously formulated approach (1.5) is a complex optimisation problem with highly non-convex and non-linear objective functions C (see Fig. 3) and E . We resort to established mathematical methods to find a solution at or near to the global optimum.

Image data, variables and involved integrals are discretised on a regular grid. The optimality requirement for the displacements and window shape variables are relaxed to the first order Karush-Kuhn-Tucker-conditions. A Levenberg-Marquardt method combined with line search, embedded into a multiscale framework, is employed for determining the step direction and scaling. No data validation algorithms (e.g. vector median filters) are applied to the variables. For details on discretisation and optimisation of the formulated problem we refer to Becker et al (2008).

3. Experiments

3.1. Synthetic Data

Our algorithm is designed to avoid smoothing over vector gradients by adapting the window shape accordingly. We verify this aim by means of a synthetic data set introduced as Case A4 in the PIV Challenge 2005. The considered region *Sinusoids I* consists of a vector field with vertical components only which vary sinus-like with wavelength from 20 to 400 pixels, see Fig. 4(a). For further details on the data set, we refer to Stanislas et al (2008).

First we applied our correlation approach but with *fixed*, round Gaussian window shapes with a radius of approximately 8 pixels. A 2-by-2-pixel-grid, six resolution levels and a scaling factor of $\sqrt{2}$ were used. The results in Fig. 4(b) show that the method is able to capture the main structures but fails to deliver accurate estimations for small wavelengths. However when additionally performing displacement estimation *combined* with window adaption, our proposed approach reconstructs the vector field with high precision even at the smallest wavelength, see Fig. 4(c). The algorithm parameter were chosen as $\sigma = 20$ and $e_{\text{outside}} = 10$.

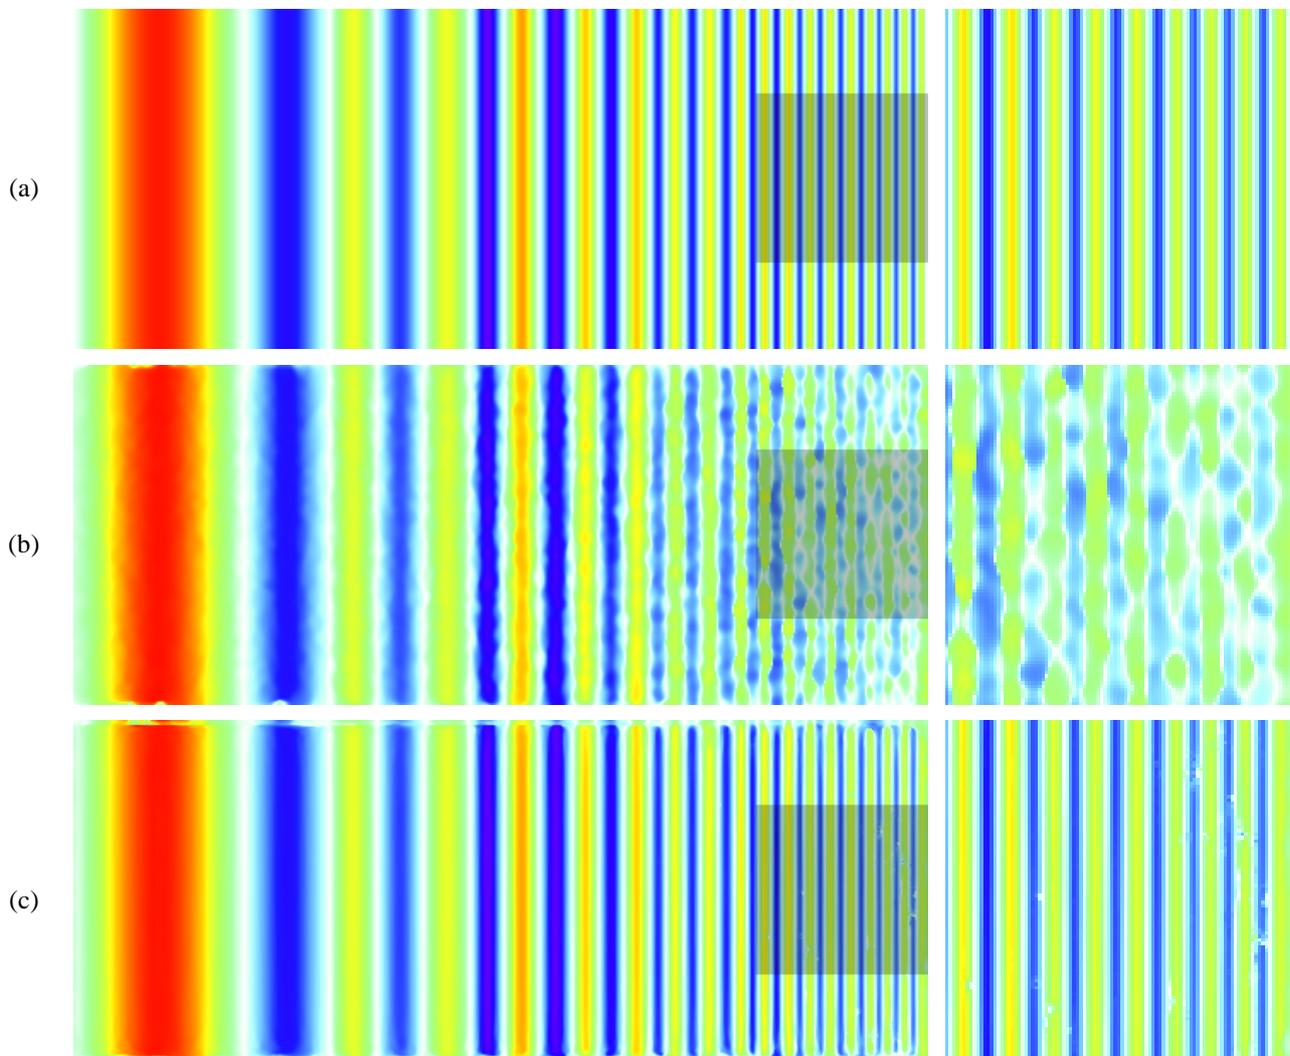


Fig. 4: Experiments with synthetic data (PIV-Challenge 2005, Case A4, Sinusoids I): Vertical component of the vector fields, (a) ground truth and shaded detail, also shown on the right; (b) correlation with fixed window shape, estimates inevitably deteriorate at small wavelengths; (c) joint correlation and window adaption can significantly improve accuracy despite spatially variant wavelengths.

3.2. Real Data

The aim of the second experiment is to evaluate whether our algorithm can improve the results of an established correlation implementation. The image data is taken from a real-world 2D-2C-PIV measurement of a wake behind a circular cylinder with a Reynolds number of 3900. Fig. 5 gives an impression on the image quality. For details on the experiment and recording of the image sequence, we refer to Carlier (2005).

Fig. 6 shows the vector field and the velocity measurement obtained from the LaVision Davis correlation software. For comparison, we further processed the data with our joint displacement and window optimisation approach. The calculation was performed on a single scale, with $\sigma=100$ and $e_{\text{outside}}=0$. Finally, the vector field was smoothed with a 3x3 second order polynomial regression filter. The result – depicted in Fig. 7 – shows how window adaption leads to smoother estimations in regions of homogenous motion, while vertexes are not affected in general shape and intensity.

Fig. 8 to Fig. 10 are detailed views of the estimated displacements and window shapes. The first illustrates the expected enlargement of window sizes in areas of homogenous motion, compared to a region of high gradients near a vortex. Fig. 9 demonstrates that window shapes do not necessarily stretch along the flow direction, but perpendicular to gradients in the vector field. In the vicinity of the vortex in Fig. 10, window sizes decrease and adapt their shape perpendicular a steep turn of the flow direction.

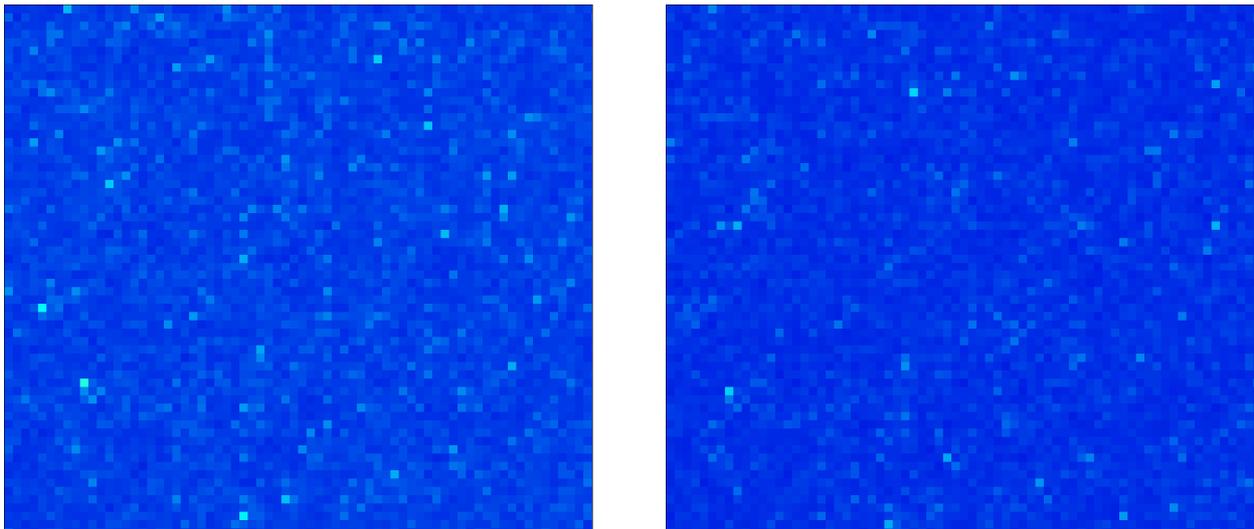


Fig. 5: Detail of the image pair recorded by a 2D-2C-PIV measurement of a wake behind a circular cylinder: Note the high noise level and small particle size.

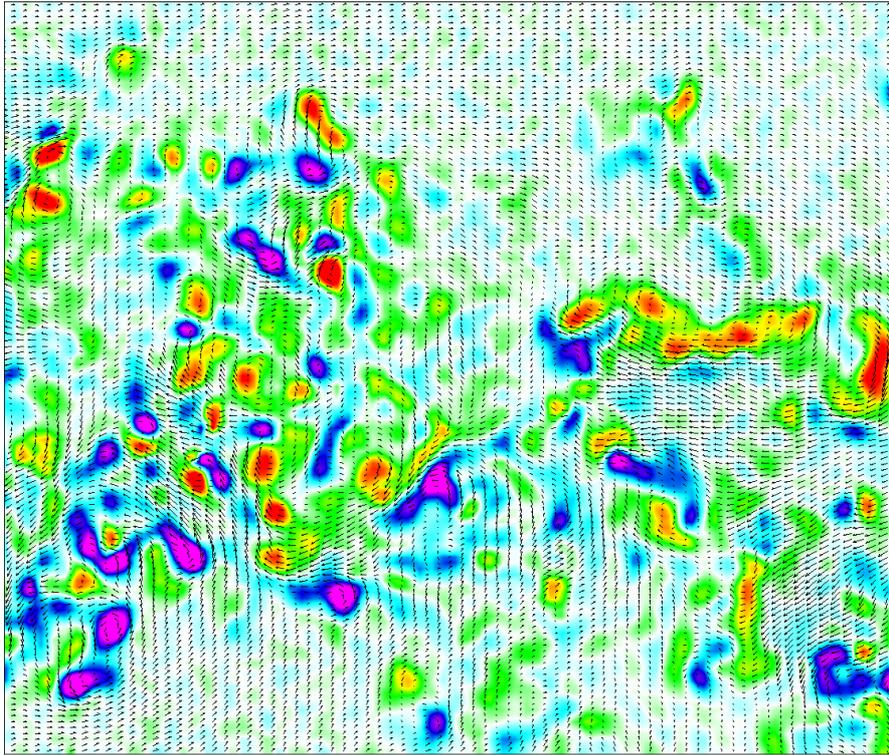


Fig. 6: Wake behind a cylinder: Velocity (after subtraction of an average displacement) and vorticity as measured by the LaVision Davis software.

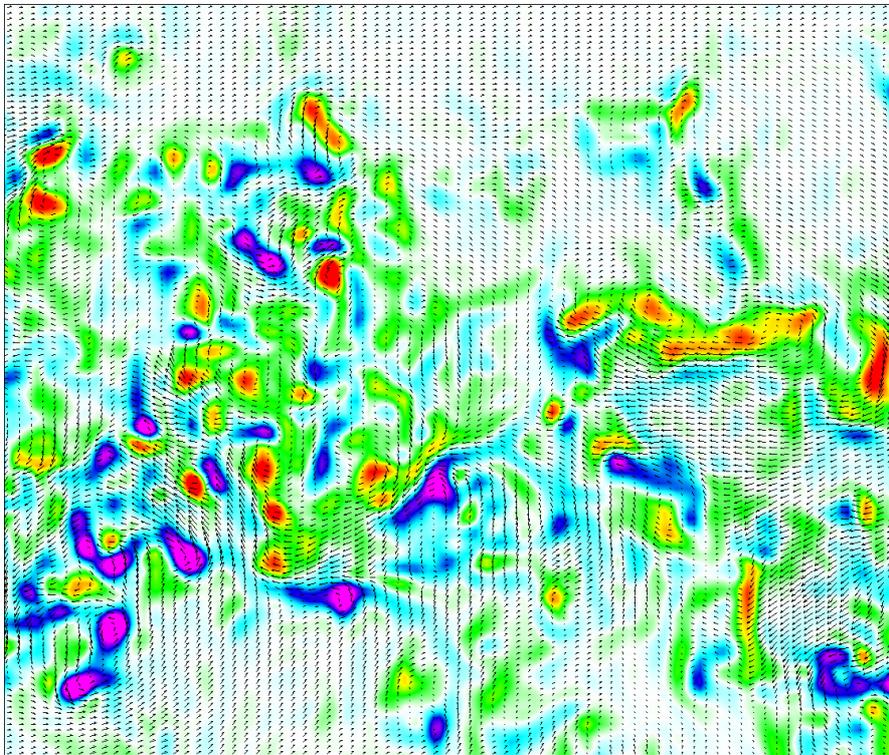


Fig. 7: Wake behind a cylinder: Velocity (after subtraction of an average displacement) and vorticity as measured by our joint displacement estimation and window adaption approach. Compared to the results in Fig. 6, regions of homogeneous motion (especially at the upper and lower border) are more clearly resolved, as correlation window sizes are enlarged there. At the same time vortex shapes are preserved as the window adaption reduces window support and adapts orientation to the gradients.

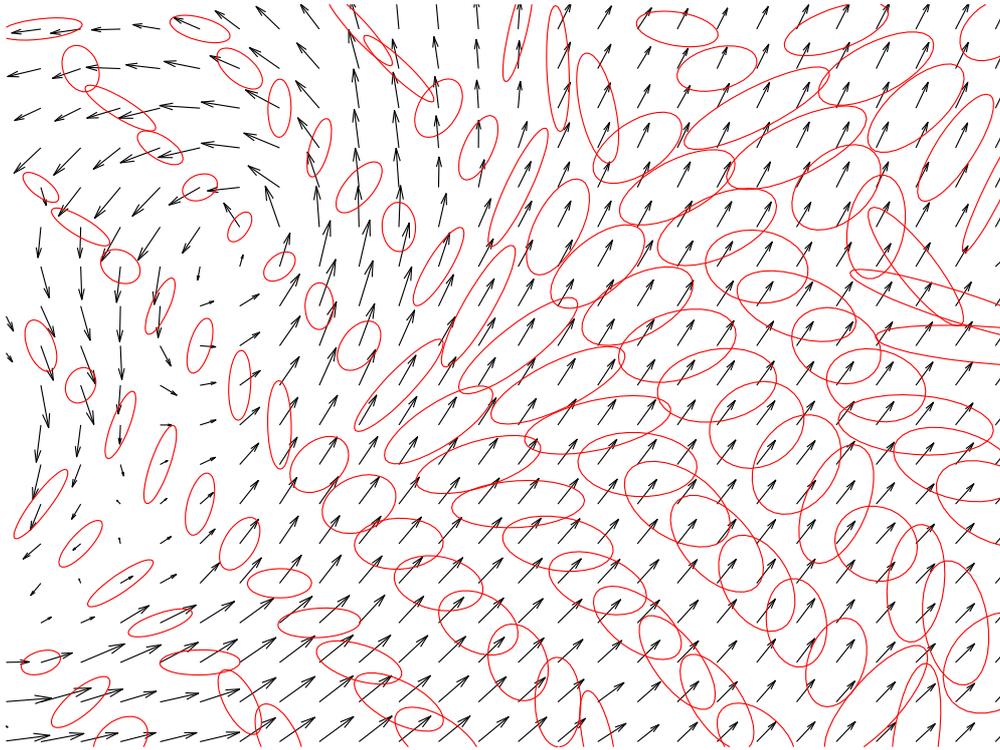


Fig. 8: Detailed view of the jointly estimated displacements (after subtracting an average displacement) and adapted correlation windows (level curve $w(x, \Sigma) = \exp(-1)$, scaled down by factor 2 for readability): Window size increases in areas of homogenous motion (middle), and decreases in presence of gradients (upper left).

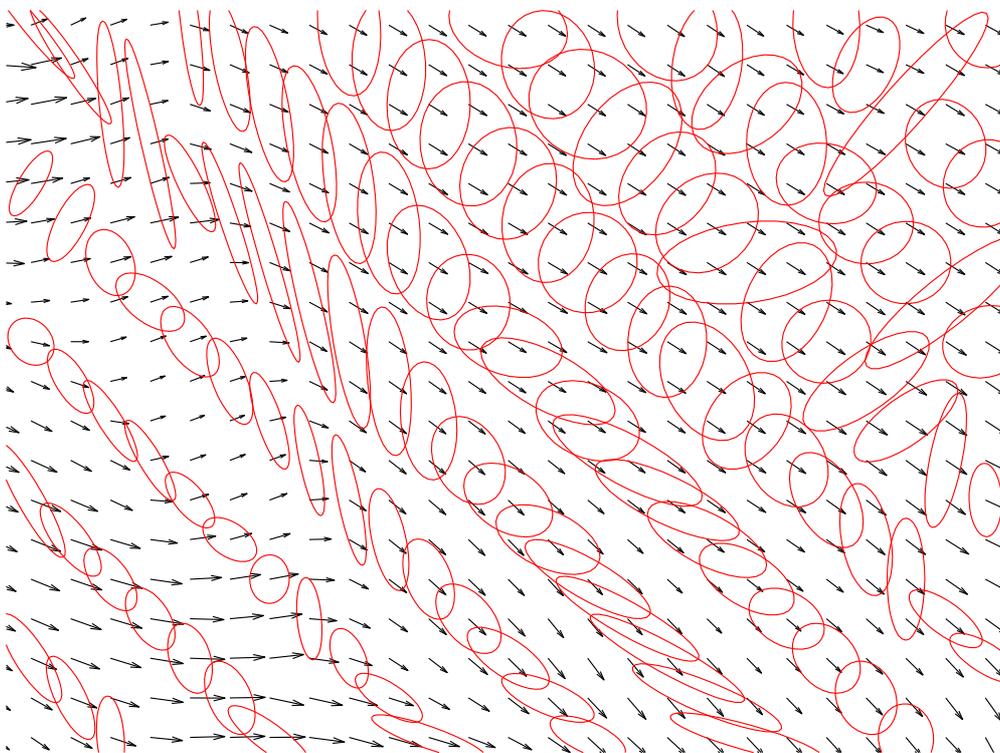


Fig. 9: Detailed view of the jointly estimated displacements and adapted correlation windows: Windows are stretched perpendicular to flow direction (upper left, along the “wave”). The window orientation does not necessarily coincide with the *direction* of the flow, but is perpendicular to its *gradient*.

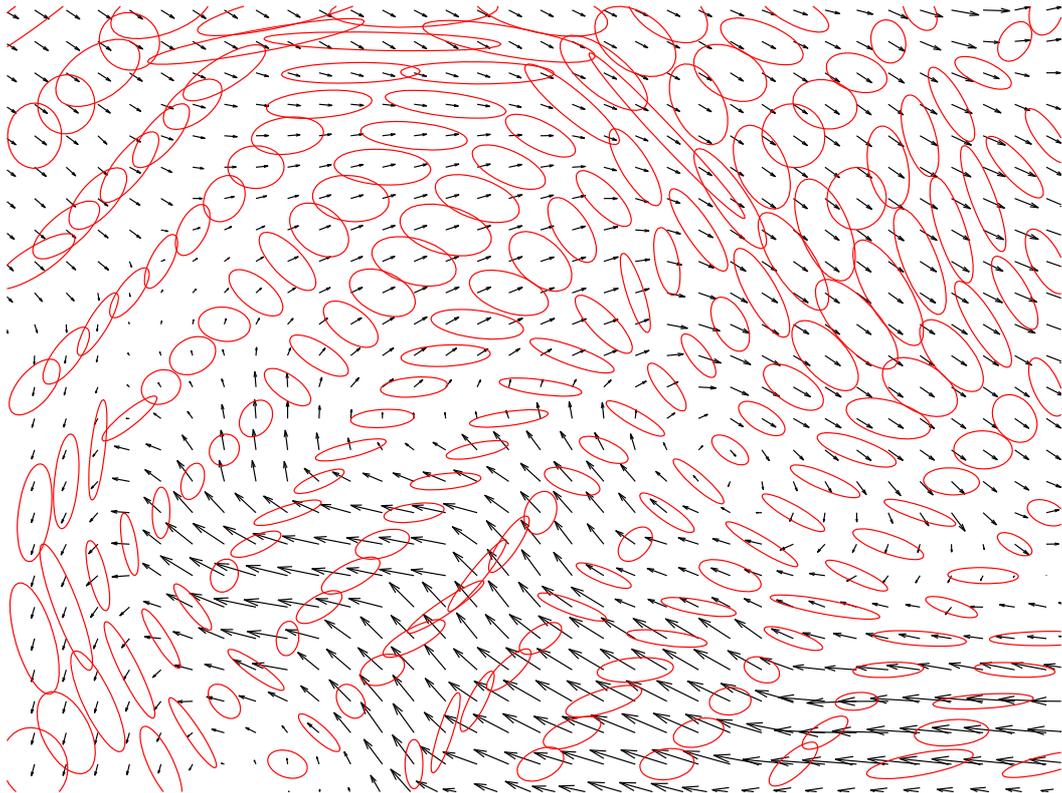


Fig. 10: Detailed view of the jointly estimated displacements and adapted correlation windows: Windows adaption around a vortex. Note that the flow is slightly “compressed” (i.e. larger gradients) along the northeast southwest axis, which also influences the window shapes accordingly.

4. Conclusion and Further Work

4.1. Conclusion

We proposed a variational approach to adaptive motion measurement for PIV data. A continuous formulation of the maximisation of the cross-correlation measure is defined using a Gaussian formed weighting function with parameterisable size, orientation and eccentricity. The displacement estimation is augmented by a sound criterion for adapting the window shape to gradients in the vector field, which directly formulates the aim to reduce the estimation error.

Experiments with synthetic data showed that our window adaption scheme can improve accuracy significantly in the presence of gradients of varying magnitude. The ability to reconstruct homogeneous regions clearly was demonstrated by means of real, turbulent data.

4.2. Further Work

The investigation of a modified, more physically motivated version of the error term is planned. In addition, choosing the parameter σ locally by means of seeding density and image noise is expected to further improve results. Further expert knowledge may be added to the error function to improve the error prediction.

Bibliography

- Becker F, Wieneke B, Yuan J, Schnörr C (2008) A Variational Approach to Adaptive Correlation for Motion Estimation in Particle Image Velocimetry. To be published in: Pattern Recognition, LNCS (2008).
- Carlier J (2005) Second set of fluid mechanic image sequences. Activity Report, European Project FLUID. Cemagref/Lavision. <http://fluid.irisa.fr/deliverables-eng.htm>
- Di Florio D, Di Felice F, Romano G P (2002) Windowing, re-shaping and re-orientation interrogation windows in particle image velocimetry for the investigation of shear flows. In: Meas Sci Technol 13:953-962
- Nogueira J, Lecuouana A, Rodríguez PA, Alfaro JA, Acosta A (2005) Limits on the resolution of correlation PIV iterative methods. Practical implementation and design of weighting functions. In: Exp Fluid 39:314-321
- Ruhnau P, Stahl A, Schnörr C (2007) Variational Estimation of Experimental Fluid Flows with Physics-based Spatio-temporal Regularization. In: Meas Sci and Technol
- Stanislas M, Okamoto K, Kähler C, Westerweel J (2008) Third International PIV-Challenge, in Exp. Fluids
- Theunissen R, Scarano F, Riethmuller ML (2007) An adaptive sampling and windowing interrogation method in PIV. Meas Sci Technol 18:275-287
- Werely ST, Meinhart CD (2001) Second-order accurate particle image velocimetry. In: Exp Fluid 31:258-268