Variational Adaptive Correlation Method for Flow Estimation

Florian Becker, Bernhard Wieneke, Stefania Petra, Andreas Schröder, and Christoph Schnörr

***Abstract*—A variational approach is presented to the estimation of turbulent fluid flow from particle image sequences in experimental fluid mechanics. The approach comprises two coupled optimisations for adapting size and shape of a Gaussian correlation window at each location and for estimating the flow, respectively. The method copes with a wide range of particle densities and image noise levels without any data-specific parameter tuning. Based on a careful implementation of a multiscale nonlinear optimisation technique we demonstrate robustness of the solution over typical experimental scenarios and highest estimation accuracy for an international benchmark dataset (PIV challenge).**

EDICS Category: ARS-IVA

I. Introduction

A. Overview

PARTICLE image velocimetry (PIV) is a non-intrusive optical measurement technique for industrial fluid flow questions [1], [2]. Small particles are added to liquids and gases and act as an indicator for the movement of the investigated substance around obstacles or in regions where flows mix. For the two-dimensional variant of PIV, a thin sheet of the volume is illuminated by a laser light, rendering the particles therein visible. A high speed camera system records an image sequence of the highlighted area. The experimental setup is sketched in Fig. 1.

The analysis of two consecutively recorded frames allows to determine the movement of particles, and in this way to measure velocity, turbulence or other derived physical properties of the fluid. Figure 2 presents an example for the image data and the derived vector field.

Cross-correlation has developed as the state of the art method for analysing PIV image data due to its robustness against distortions typical for this application [3]. These are high image noise levels caused by high shutter speeds, as well as varying laser output and particles leaving or entering the illuminated plane.

In this work we formulate a variational approach to velocity measurement based on cross-correlation. Furthermore, we replace the usually employed sharp square correlation window by a Gaussian weighting function. A mathematically sound and novel criterion to adapt window sizes and shapes is proposed, which directly formulates the aim to improve the accuracy of the velocity estimation, especially in the presence of vortices as demonstrated in Fig. 3.

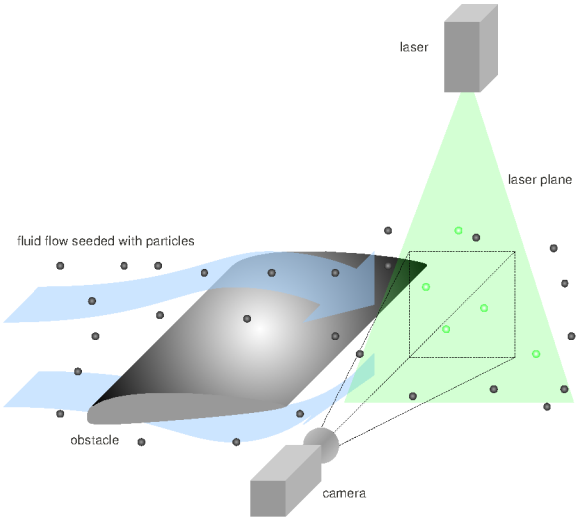

Fig. 1. Particle Image Velocimetry (PIV) experimental setup: The flow of a fluid behind an obstacle shall be examined. To this end, it is seeded with small particles. A laser projects a thin plane into the region of interest and illuminates the particles therein. Fluid motion can be measured by analysing the image sequence recorded by a camera system.

Both displacement estimation and window adaptation are joint into a coupled optimisation problem and solved using methods for non-linear and non-convex optimisation.

We demonstrate the ability of the approach to significantly improve the accuracy of flow measurements with synthetic examples as well as real turbulent data. For an international benchmark data set our method outperforms most of the concurring approaches.

B. Related Work and Contribution

A vast body of literature exists on all aspects of the application of cross-correlation for analysing PIV image data, here we only refer to [1] for an excellent overview. In a typical implementation, an exhaustive search over integer displacements is performed to find the highest correlation peak which corresponds to the most probable displacement of the considered region between two consecutive frames. This process can be speeded up using Fast Fourier Transformation. The correlation function is then interpolated to gain sub-pixel accuracy. In contrast, we present a *variational* approach which determines the displacements by *continuously* maximising the cross-correlation measurement.

F. Becker, S. Petra and C. Schnörr are with the Heidelberg Collaboratory for Image Processing and the Image & Pattern Analysis group, University of Heidelberg, Germany, e-mail: {becker,petra,schnoerr}@math.uni-heidelberg.de; B. Wieneke is with Lavision GmbH, Germany, e-mail: bwieneke@lavision.de; A. Schröder is with the Institute of Aerodynamics and Flow Technology, German Aerospace Center, e-mail: andreas.schroeder@dlr.de

Manuscript received xxxxx xx, xxxx; revised xxxxx xx, xxxx

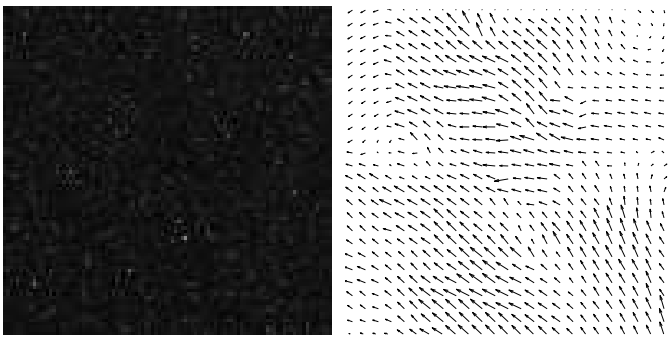

(a) recorded frame (b) reconstructed displacement field

Fig. 2. Particle Image Velocimetry (PIV): **(a)** Detail of an image frame recorded in a real fluid experiment (see Sect. IV-D), showing particles (polyglycol diluted in water) in an air flow. The image is $100\,\text{px} \times 100\,\text{px}$ in size, corresponding to $7.5\,\text{mm} \times 7.5\,\text{mm}$ in the illuminated plane. The particles have a diameter of less than $10\,\mu\text{m}$ i.e. less than a pixel in the projection, which prevents the application of standard differential optical flow techniques. Image noise level is high. **(b)** Example of a reconstructed displacement field. Each arrow describes the motion vector estimated at its origin.

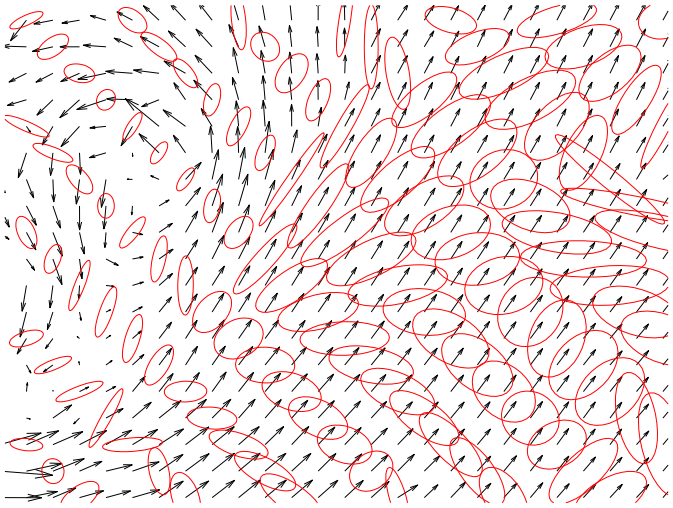

Fig. 3. Estimated displacement field in a real turbulent PIV experiment (arrows) and some of the adapted correlation windows (represented by ellipse-shaped level contours). For each measurement, our approach balances the error caused by a high amount of image noise compared to the supporting region, and violation of the assumption, that the flow is constant within the window, such as in the upper left region, where high gradients dominate. In contrast, windows can be chosen larger in homogeneous regions (right). This spatially-dependent adaptivity of windows emerges from our variational approach.

Variational approaches based on the linearised optical flow constraint, originating from the work of Horn and Schunck [4], have also been applied to measure motion in PIV [5]. The dense vector field representation allows to incorporate prior knowledge on the vector field, such as incompressibility, see e.g. [6], [7]. However, the implied brightness constancy assumption often does not hold in real PIV data which is taken into account in [8] by modifying the data term.

In [9], cross-correlation is considered as one of three similarity measurements between image pairs. Based on this and additional spatial regularisation, a variational approach for image registration is formulated and solved using partial differential equations.

The two-scale approach in [10] combines the advantages of optical flow based methods and cross-correlation: an optical flow approach with physically sound regularisation terms, which penalise large variations in the rotation and divergence of the flow, is endowed with an additional data term. Similarity to a coarse vector field originating from a local correlation approach is enforced. The displacement field estimated using cross-correlation is used to *initialise* a variational optical flow approach in [11]. For a comprehensive synopsis on variational methods for fluid flow measurement we refer to [12].

Much effort has been put into improving the spatial resolution of cross-correlation methods [13], [14] by replacing the fixed square interrogation windows by appropriate alternatives. The authors of [15] investigate a class of cone-shaped weighting functions and optimise the shape parameters by means of the frequency-response, however not with respect to a specific image data set. In [16], [17] the size of square windows is locally adapted to the signal quality and spatial fluctuations in the flow. Window adaptation is used in [18] at interfaces to fixed objects in the scene. The authors of [19] propose a criterion based on the flow gradients and image quality to select the optimal shape of an elliptical window. In [20] a Gaussian weighting function is stretched and rotated along the measured mean displacement. Gaussian weights are used both in a local [21] and global context [22] for smoothing the optical flow constraint, however with isotropic windows of fixed size common for all positions. In our work, the correlation window is also described by a "soft" Gaussian weighting function. However, we formulate a sound criterion for the location-dependent choice of the window shape parameters by means of an error model function. The window adaptation consists of finding the window shape which minimises the predicted measurement error.

Our contribution is a variational formulation for a correlation-based approach for measuring motion in PIV image pairs. A Gaussian weighting function controls the image region considered in the displacement estimation. The shape of the window is controlled by means of a function which approximates the expected measurement error. Minimisation leads to the optimal window shape with respect to this error model. Displacement measurement and window adaptation are formulated as a pair of interdependent optimisation problems. It is solved via a multiscale gradient-based algorithm.

This work summarises and extends the results in [23]. An abridged version was published in [24] with the focus on image processing. In [25] we investigate our approach from the applied fluid mechanics point of view.

C. Organisation

In Sect. II we formulate our approach to adaptive fluid flow measurement as a continuous optimisation problem. Section III details on the discretisation and the employed optimisation method. We verify both in the experimental section (Sect. IV) and conclude in Sect. V.

II. Approach

A. Problem Definition

Given a pair of images, g_1, g_2, defined on the image domain $\Omega \subset \mathbb{R}^2$, we are interested in a vector field $\mathbf{u} : \Omega \mapsto \mathbb{R}^2$

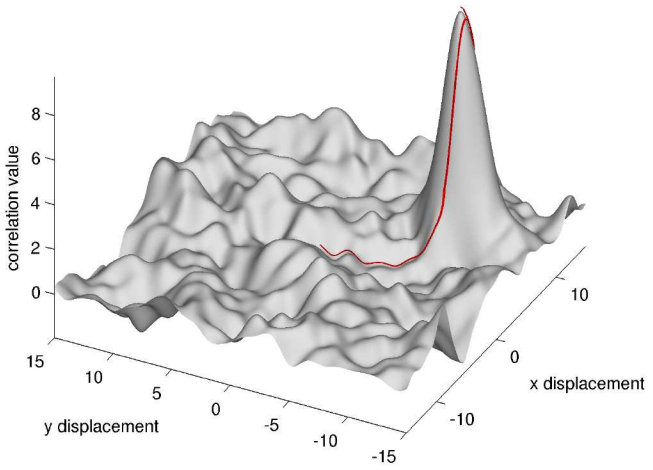

Fig. 4. Value of the cross-correlation-function, $-C(u, \Sigma, x)$ between two particle images and a path (starting in $(0, 0)$) of a gradient-based optimisation approach. Multiscale techniques are employed to avoid that the process stops in a local optimum and to find the correlation peak in $(8, -8)$.

which locally describes the displacement of image structures from g_1 to g_2.

B. Variational Correlation

For the estimation of the displacement at $x \in \Omega$, we consider the following optimisation problem:

$$\min_{u \in \mathbb{R}^2} C(u, \Sigma, x) \;, \quad \text{with} \tag{1}$$

$$C(u, \Sigma, x) = -\int_{\mathbb{R}^2} w(y - x, \Sigma)\, g_1\Big(y - \frac{u}{2}\Big)\, g_2\Big(y + \frac{u}{2}\Big)\, \mathrm{d}y \tag{2}$$

Basically, the considered regions in g_1 and g_2 are shifted until they fit best with respect to the correlation measurement. The function $w(x, \Sigma) \in [0, 1]$ weights the impact of the image region in the vicinity of x. In Sect. II-C we provide further details and propose a method to adapt the parameter Σ which controls shape and size with the aim to improve accuracy.

The objective function $C(u, \Sigma, x)$ is non-linear and highly non-convex in u as Fig. 4 illustrates. Thus, the employed optimisation method has to circumnavigate several local optima to reach the correlation peak. We will come back to this issue in Sect. III-B.

Finally, we extend the problem definition to estimating a complete flow field on Ω and define the variable domains

$$\mathcal{U} := \left\{ \mathbf{u} : \Omega \mapsto \mathbb{R}^2 \,\middle|\, \int_\Omega \|\mathbf{u}(x)\|_2^2 \,\mathrm{d}x < \infty \right\} \tag{3}$$

$$\text{and } \mathcal{S} := \left\{ \mathbf{\Sigma} : \Omega \mapsto S \,\middle|\, \int_\Omega \|\mathbf{\Sigma}(x)\|_\mathrm{F}^2 \,\mathrm{d}x < \infty \right\} . \tag{4}$$

The set $S \subset \mathcal{S}_{++}^2$ denotes the allowed window shapes as detailed further below, where $\mathcal{S}_{++}^2$ is the set of symmetric, positive 2×2 matrices. $\|\cdot\|_\mathrm{F}$ denotes the Frobenius matrix norm. For given, fixed window shapes $\mathbf{\Sigma}$, we determine a flow field $\mathbf{u}$ on Ω by solving the following optimisation problem:

$$\min_{\mathbf{u} \in \mathcal{U}} C(\mathbf{u}, \mathbf{\Sigma}) \tag{5}$$

$$\text{with } C(\mathbf{u}, \mathbf{\Sigma}) := \int_\Omega C(\mathbf{u}(x), \mathbf{\Sigma}(x), x)\, \mathrm{d}x \tag{6}$$

Fig. 5. Some possible shapes of the Gaussian weighting function as grey-value plot (white=0, black=1), from top to bottom row: varying radius, anisotropy and orientation

C. Window Adaptation

Instead of a square or rectangular window, we choose the weighting function as a non-normalised Gaussian function:

$$w(x, \Sigma) := \frac{\mathrm{G}(\Sigma, x)}{\mathrm{G}(\Sigma, 0)} = \exp\left(-\frac{1}{2} x^\top \Sigma^{-1} x\right) , \tag{7}$$

with $\mathrm{G}(\Sigma, x) = \frac{1}{2\pi\sqrt{\det \Sigma}} \exp\big(-\frac{1}{2} x^\top \Sigma^{-1} x\big)$ being the normalised two-dimensional Gaussian function with covariance matrix $\Sigma \in \mathcal{S}_{++}^2$, centred at $x = 0$. The shape parameter Σ allows to select the image region considered for the correlation measurement. It provides high flexibility in three degrees of freedom, which can be interpreted as geometric properties of the ellipsoidal contour line $\big\{ x \in \mathbb{R}^2 \,\big|\, w(x, \Sigma) = \exp(-1) \big\}$, namely semi-minor radius $r \geq 0$, anisotropy $a \in [0, 1)$ and orientation $\alpha \in [0, \pi)$:

$$\Sigma(r, a, \alpha) = \frac{1}{2} Q D Q^\top \tag{8}$$

$$\text{with } Q = \begin{pmatrix} \cos\alpha & \sin\alpha \\ -\sin\alpha & \cos\alpha \end{pmatrix} \text{ and } D = \operatorname{diag}\left(\frac{r^2}{1 - a^2}, r^2\right)$$

Figure 5 illustrates the influence of the three parameters. However, in the following we parametrise the window shape by the matrix representation Σ.

The choice of the window shape is motivated by the following consideration: when cross-correlation is employed for estimating displacements, it is implicitly assumed that motion within the considered image regions is uniform. This, however, only holds true in very simple cases and leads to large estimation errors in regions with motion gradients as they occur in typical fluid flows. Decreasing the window size reduces this effect, however at the costs of a smaller supporting area, a smaller number of observed particles and thus a higher impact of image noise.

With the improvement of measuring accuracy in mind, we model the described effects and find a trade-off between them by solving the following optimisation problem for each location $x \in \Omega$

$$\mathbf{\Sigma}(x) \in \arg\min_{\Sigma \in S} E(\Sigma, \mathbf{u}, x)$$
$$\text{with } E(\Sigma, \mathbf{u}, x) := E_\mathrm{homog}(\Sigma, \mathbf{u}, x) + E_\mathrm{noise}(\Sigma) \;, \tag{9}$$

assuming that a displacement field $\mathbf{u}$ is given. Furthermore, we incorporate lower (λ_min) and upper (λ_max) bounds

on the window size by defining the constraint set $S := \left\{ \Sigma \in \mathcal{S}^2_{++} \,\middle|\, \lambda_{\min} I \preceq \Sigma \preceq \lambda_{\max} I \right\}$. I denotes the identity matrix of appropriate size and $S \preceq T$ indicates positive semi-definiteness of $T - S$ for symmetric matrices S, T.

The error model function $E(\Sigma, \mathbf{u}, x)$ is composed of the following two terms:

a) Homogeneity Term: The first part of the objective (9) describes the error caused by the violation of the assumption that the observed motion within the chosen window is homogeneous:

$$E_{\text{homog}}(\Sigma, \mathbf{u}, x) := \int_{\mathbb{R}^2} w(y - x, \Sigma)\, e(x, y, \mathbf{u})\, \mathrm{d}y \tag{10}$$

$$e(x, y, \mathbf{u}) := \begin{cases} \|\mathbf{u}(y) - \mathbf{u}(x)\|_2^2 & \text{if } y \in \Omega \\ e^2_{\text{outside}} & \text{otherwise} \end{cases} \tag{11}$$

The function $e(x, y, \mathbf{u})$ measures the squared Euclidean distance of $\mathbf{u}(y)$ to the displacement in the point of interest, $\mathbf{u}(x)$, while a constant error value is assumed for using regions outside the image domain Ω. The errors are weighted by the window function which is parameterised by Σ.

b) Noise Term: The second term in (9) describes the impact of image sensor noise and unpaired particles on the accuracy. We define it as

$$E_{\text{noise}}(\Sigma) = \frac{\sigma^2}{2\pi\sqrt{\det \Sigma}}\,, \tag{12}$$

where σ is a parameter which describes the image noise level. Intuitively spoken, this term describes the expectation that the error reduces when the measurement support, $\int w(x, \Sigma)\, \mathrm{d}x = 2\pi\sqrt{\det \Sigma}$ increases for larger windows Σ. A more detailed derivation of this term can be found in the Appendix.

c) Global Window Adaptation: Finally, we extend the local window estimation to

$$\min_{\mathbf{\Sigma} \in \mathcal{S}} E(\mathbf{\Sigma}, \mathbf{u}) \quad \text{with } E(\mathbf{\Sigma}, \mathbf{u}) := \int_\Omega E(\mathbf{\Sigma}(x), \mathbf{u}, x)\, \mathrm{d}x\,,$$

which optimises the window shapes globally in terms of the matrix-valued function $\mathbf{\Sigma} \in \mathcal{S}$.

D. Joint Approach

In Sect. II-B we introduced a motion estimation approach and presumed that the window parameters are given. In contrast, in Sect. II-C we fixed a displacement field and adapted the correlation windows to it. We describe this chicken-and-egg-dependency as a mathematically tractable problem:

$$\mathbf{u}^* \in \arg\min_{\mathbf{u} \in \mathcal{U}} C(\mathbf{u}, \mathbf{\Sigma}^*) \;\text{ and }\; \mathbf{\Sigma}^* \in \arg\min_{\mathbf{\Sigma} \in \mathcal{S}} E(\mathbf{\Sigma}, \mathbf{u}^*) \tag{13}$$

The two optimisation problems have non-linear and non-convex objective functions each, are interconnected through the variables $\mathbf{\Sigma}^*$ and $\mathbf{u}^*$, and thus have to be solved jointly.

Note, that $C(\mathbf{u}, \mathbf{\Sigma}^*)$ is only minimised with respect to $\mathbf{u}$, but *not* $\mathbf{\Sigma}^*$, as the window shapes should only be steered by the error model function and not by the correlation measurement. If image data should be considered in the window choice, an additional term should be incorporated into $E(\mathbf{\Sigma}, \mathbf{u}^*)$.

III. Discretisation and Optimisation

A number of carefully chosen approximations and relaxations were applied to make the optimisation problem (13) tractable.

A. Discretisation

The functions $\mathbf{u}$ and $\mathbf{\Sigma}$ are discretised component-wise on a regular grid X_{V} with spacing a_{V} at coordinates $x_i \in X_{\text{V}}$. Furthermore, we define $\mathbf{u}_i := \mathbf{u}(x_i)$ and $\mathbf{\Sigma}_i := \mathbf{\Sigma}(x_i)$. Using finite elements with piecewise linear basis functions $\varphi_i(x)$, we approximate the functions as

$$\mathbf{u}(x) \approx \sum_{x_i \in X_{\text{V}}} \varphi_i(x) \mathbf{u}_i \quad \text{and} \quad \mathbf{\Sigma}(x) \approx \sum_{x_i \in X_{\text{V}}} \varphi_i(x) \mathbf{\Sigma}_i\,.$$

Note that it is possible to extend the method to arbitrary grids, e.g. irregular ones that adapt to the seeding density, as it is used in [26]. The integrals in C and E are also discretised using the introduced finite elements:

$$C(\mathbf{u}, \mathbf{\Sigma}) \approx \sum_{x_i \in X_{\text{V}}} \int_\Omega \varphi_i(x)\, \mathrm{d}x\, C(\mathbf{u}_i, \mathbf{\Sigma}_i, x_i)$$

$$E(\mathbf{\Sigma}, \mathbf{u}) \approx \sum_{x_i \in X_{\text{V}}} \int_\Omega \varphi_i(x)\, \mathrm{d}x\, E(\mathbf{\Sigma}_i, \mathbf{u}, x_i)$$

Note that $A_i := \int_\Omega \varphi_i(x)\, \mathrm{d}x$ evaluates to a_V^2 almost everywhere. The nested integrals in $C(u, \Sigma, x)$ and $E_{\text{homog}}(\Sigma, \mathbf{u}, x)$ are discretised accordingly.

For windows of reasonable size the function $w(x, \Sigma)$ incorporated in both terms weights only few terms with considerable impact. In order to reduce computational effort, we limit evaluation to a bounding box which contains all $y \in \Omega$, such that $w(y - x, \Sigma) \leq 10^{-3}$.

The image data g_1, g_2 is given on a regular grid with a spacing typically smaller than a_{V}. They are transferred into a cubic spline representation with all values outside the image domain defined to be zero. Using an efficient implementation based on [27], it is possible to evaluate the function value g_i, its gradient ∇g_i and second derivatives $\mathrm{H}\, g_i$.

B. Optimisation

1) Barrier Function: The constraints $\Sigma \in S$ are incorporated into the energy function using logarithmic barriers,

$$B_S(\Sigma) := -\mu \left(\log\det(\Sigma - \lambda_{\min}) + \log\det(\lambda_{\max} I - \Sigma) \right).$$

The penalty weight is $\mu := 10^{-2}$ throughout the work. Then we minimise $E_S(\Sigma, \mathbf{u}, x) := E(\Sigma, \mathbf{u}, x) + B_S(\Sigma)$ instead of (9), which is – up to the symmetry of Σ – an unconstrained problem.

2) Single Scale Optimisation: A major simplification of the problem is to replace both minimality objectives by the stationary conditions

$$\nabla_{\mathbf{u}_i} C(\mathbf{u}, \mathbf{\Sigma}) = 0 \qquad \forall x_i \in X_{\text{V}} \tag{14a}$$

$$\text{and } \nabla_{\mathbf{\Sigma}_i} E(\mathbf{\Sigma}, \mathbf{u}) = 0 \qquad \forall x_i \in X_{\text{V}}\,. \tag{14b}$$

A Newton step with respect to all displacement and window shape variables is employed to find a set of $\mathbf{\Sigma} \in \mathcal{S}$ and $\mathbf{u} \in \mathcal{U}$

that satisfy these conditions. It is extended by a line search method to avoid local maxima and saddle points.

Note, that although each equality constraint in (14a)-(14b) is a nonlinear and non-convex function in both $\mathbf{u}$ and $\mathbf{\Sigma}$, they strongly simplify to

$$\begin{aligned} \nabla_{\mathbf{u}_i} C(\mathbf{u},\mathbf{\Sigma}) &= \nabla_{\mathbf{u}_i} A_i C(\mathbf{u}_i,\mathbf{\Sigma}_i,x_i) = 0 \qquad \forall x_i \in X_\mathrm{V} \\ \text{and } \nabla_{\mathbf{\Sigma}_i} E(\mathbf{\Sigma},\mathbf{u}) &= \nabla_{\mathbf{\Sigma}_i} A_i E(\mathbf{\Sigma}_i,\mathbf{u},x_i) = 0 \qquad \forall x_i \in X_\mathrm{V}\ . \end{aligned}$$

Thus, the displacements can be updated independently of each other which is a consequence of the fact that we did not add a spatial regularisation term on $\mathbf{u}$. In the same way, the refinement of the window shape parameters is independent in the coordinates. The optimisation loop can be summarised as:

procedure SINGLESCALESOLUTION($\mathbf{u}^{(1)}$,$\mathbf{\Sigma}^{(1)}$)
  $k \leftarrow 1$
  repeat
    for all $x_i \in X_V$ **do**
      $\mathbf{u}_i^{(k+1)} \leftarrow$ VARIABLEUPDATE($C(\mathbf{u},\mathbf{\Sigma}), \mathbf{u}_i, (\mathbf{u}^{(k)}, \mathbf{\Sigma}^{(k)})$)
      $\mathbf{\Sigma}_i^{(k+1)} \leftarrow$ VARIABLEUPDATE($E(\mathbf{\Sigma},\mathbf{u}), \mathbf{\Sigma}_i, (\mathbf{\Sigma}^{(k)}, \mathbf{u}^{(k)})$)
    end for
    $k \leftarrow k+1$
  until stopping criterion fulfilled
  return $(\mathbf{\Sigma}^{(k)}, \mathbf{u}^{(k)})$
end procedure

An upper bound on the change of the variables is used as stopping criterion. The function VARIABLEUPDATE improves the solution x_0 with respect to f by updating a subset of variables y.

procedure VARIABLEUPDATE(f,y,x_0)
  $g \leftarrow \nabla_y f(x_0)$, $H \leftarrow \mathrm{H}_y f(x_0)$ ▷ gradient, Hessian
  $\Delta y \leftarrow -(H + \lambda I)^{-1} g$ ▷ Newton step direction w.r.t. y
  $\alpha \leftarrow \alpha_{\max}$, $x_\alpha \leftarrow x_0$
  while $\alpha \geq \alpha_{\min}$ **do** ▷ line search
    $x_\alpha|_y \leftarrow x_0|_y + \alpha \Delta y$ ▷ update only variables y
    if $f(x_\alpha) < f(x_0)$ **then**
      return x_α ▷ update successful
    end if
    $\alpha \leftarrow \beta\alpha$
  end while
  return x_0 ▷ update failed
end procedure

The parameters were chosen conservatively: $\alpha_{\min} = 10^{-9}$, $\alpha_{\max} = 1 - 10^{-3}$, $\beta = 10^{-1}$ and $\lambda = 100$.

3) Multiscale Optimisation: As indicated by Fig. 4, the problem has many local minima which we intend to circumnavigate by wrapping a multiscale framework around the optimisation loop. To this end, we represent the problem at the original as well as a couple of coarser resolutions. The grid spacings (data and variable) enlarge by factor $s > 1$ when descending one level. E.g. for a 5-level dyadic pyramid, we have $s = 2$ and we denote the resolution scales – from finest to coarsest – as $\{1, 2, 4, 8, 16\}$.

The multiscale framework first recursively transfers image and initial variable values from the finest to the coarsest level. Then at each resolution the estimated solution of the next coarser level act as initialisation for the variable refinement in SINGLESCALESOLUTION.

Displacement variables are re-sampled to finer or coarser grids using cubic spline interpolation. A small binomial low-pass filter is used to avoid aliasing while down-sampling. The multiscale image representation is created with the same technique. The re-sampling process of the window shape parameter is slightly more complex, as the constraint $\mathbf{\Sigma} \in S$ has to be conserved. However, simple component-wise bi-linear interpolation guarantees that the re-sampled value lies in the convex hull of the interpolated values, and thus in S. The same argument holds for applying low-pass filters as long as their coefficients add up to one, such as it is the case for the employed binomial filters before down-sampling. Given the image data and (a possibly zero) initial solution, the overall optimisation can be summarised as:

procedure MULTISCALESOLUTION($g_1^{[1]}$,$g_2^{[1]}$,$\mathbf{u}^{[1]}$,$\mathbf{\Sigma}^{[1]}$)
  for $l = 2, 3, \ldots, l_{\max}$ **do** ▷ fine to coarse
    create $g_i^{[l]}$ by downsampling $g_i^{[l-1]}$, for $i \in \{1, 2\}$
    create $\mathbf{u}^{[l]}$ and $\mathbf{\Sigma}^{[l]}$ by down-sampling $\mathbf{u}^{[l-1]}$ and $\mathbf{\Sigma}^{[l-1]}$
  end for
  for $l = l_{\max}, l_{\max} - 1, \ldots, 2$ **do** ▷ coarse to fine
    $(\mathbf{u}^{[l]}, \mathbf{\Sigma}^{[l]}) \leftarrow$ SINGLESCALESOLUTION($\mathbf{u}^{[l]}, \mathbf{\Sigma}^{[l]}$)
    create $\mathbf{u}^{[l-1]}$ and $\mathbf{\Sigma}^{[l-1]}$ by up-sampling $\mathbf{u}^{[l]}$ and $\mathbf{\Sigma}^{[l]}$
  end for
  $(\mathbf{u}^{[1]}, \mathbf{\Sigma}^{[1]}) \leftarrow$ SINGLESCALESOLUTION($\mathbf{u}^{[1]}, \mathbf{\Sigma}^{[1]}$)
  return ($\mathbf{u}^{[1]}$,$\mathbf{\Sigma}^{[1]}$)
end procedure

Further details on the implementation can be found in [23].

IV. EXPERIMENTS

In our experiments we investigated the basic properties of the window adaptation (Sect. IV-A, IV-B), and evaluated the joint approach with synthetic benchmark data (Sect. IV-C) as well as real-world data (Sect. IV-D).

The proposed methods was implemented mostly in MATLAB. Geometric properties of a window Σ, such as radius, refer to the level contour $\{x \in \mathbb{R}^2 | w(x, \Sigma) = \exp(-1)\}$, which is also used for visualisation. No additional displacement filters (e.g. vector median) were applied.

A. Window Adaptation Strategies

The following experiments were designed to estimate the potential of the proposed error model to improve the accuracy of the variational correlation method. As we want to concentrate on the suitability of error model function and not on errors caused by the continuous optimisation process, we simplify the method as follows: in order to avoid sub-optimal local minima, the optimisation of the correlation is initialised by the ground truth displacement. For the same reason, we do not adapt the window continuously but evaluate the deviation from ground truth for 975 window shapes using (8) with varying radius r, orientation α and anisotropy a,

$$S := \left\{ \Sigma(r,a,\alpha) \;\middle|\; \begin{array}{l} r \in \left\{ 2^{\frac{i}{2}} \,\middle|\, i \in \{-4,-3,\ldots,10\} \right\}, \\ a \in \left\{ 1 - 2^{\frac{i}{4}} \,\middle|\, i \in \{0,1,\ldots,8\} \right\}, \\ \alpha \in \left\{ \frac{i}{8}\pi \,\middle|\, i \in \{0,1,\ldots,7\} \right\} \end{array} \right\}.$$

(a) vertical displacement component (b) detail

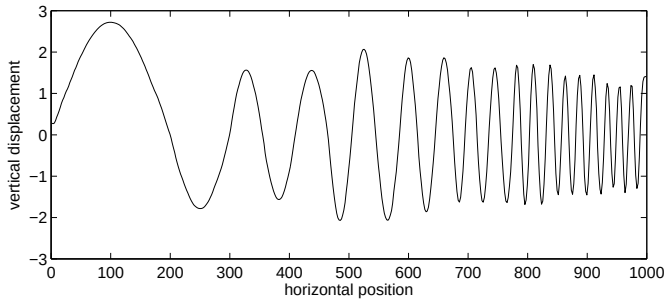

(c) profile of the vertical displacement component

Fig. 6. Synthetic vector field, *PIV-Challenge 2005*, *Case A4*: **(a)** vertical displacement component (-3 px (black) to +3 px (white) upwards), a detailed view of the highlighted region is given in **(b)**; **(c)** plot of the vertical component of the displacements common for both data sets, *Sinusoids I* and *Sinusoids II*. The vector field has a zero horizontal component, while the vertical component is piecewise described by sine functions with decreasing wavelength (400 to 20 px) and varying amplitude (around 2 px). While the first data set contains no image distortions, 3% pixel noise and 20% unpaired particles were added to the second one.

The correlation method was applied to 160 locations $X \subset \Omega$ in two synthetic data sets with motion gradients of varying degree. Both have a common motion pattern, which is illustrated in Fig. 6, but differ in the amount of image distortions. The deviation of the estimated displacement from ground truth at position $x \in X$ using the (fixed) window shape $\Sigma \in S$ is then measured by their Euclidean distance and is denoted by $\varepsilon(x, \Sigma)$.

We considered our adaptation approach as one of three strategies (indexed by $i \in \{1, 2, 3\}$) which choose the window shape $\Sigma_i(x)$ from the set S for a position x. Their performance is compared on the basis of the mean error $\mu_i := \frac{1}{|X|} \sum_{x \in X} \varepsilon(x, \Sigma_i(x))$. The three strategies are in detail:

Strategy 1 (Oracle): This hypothetical strategy "magically" knows the values $\varepsilon(x, \Sigma)$ a priori and can always choose the optimal window:

$$\Sigma_1(x) := \arg\min_{\Sigma \in S} \varepsilon(x, \Sigma)$$

Thus, the mean error of this strategy, denoted as μ_1^*, provides a lower bound for all strategies under this conditions.

Strategy 2 (Error Model): The second strategy represents the proposed window adaptation method. For each position, the window shape is chosen such that it is optimal with respect to the defined error model function:

$$\Sigma_2(x) := \arg\min_{\Sigma \in S} E(\Sigma, \mathbf{u}, x)$$

This optimisation problem is solved by enumerating all elements of S. Furthermore, the ground truth vector field is used for $\mathbf{u}$ to exclude influences caused by inaccuracies in the displacement estimation. The mean error is denoted as $\mu_2(\sigma)$ and depends on the choice of σ, while the parameter e_{outside} was set to zero.

Strategy 3 (Fixed Radius): A naive strategy is to choose the window radius r a priori and uniformly for all position.

$$\Sigma_3(x) := \Sigma(r, 0, 0)$$

The associated quality measurement is denoted by $\mu_3(r)$.

For the latter two strategies we also define σ^* and r^* which minimise the corresponding mean errors and define these values as $\mu_2^* := \mu_2(\sigma^*)$ and $\mu_3^* := \mu_3(r^*)$, respectively.

Figure 7 visualises and lists the results. For both data sets, error rate improves by approximately 50% compared to the fixed-radius strategy if window adaptation based on the proposed error model is used. Furthermore, the mean error μ is less sensible to the choice of the parameter σ than to the window radius r.

B. Synthetic Vector Fields

For arbitrary vector fields, the optimal window shape with respect to the proposed error model can form complex structures. To demonstrate the behaviour of the window adaptation methods, we investigate a couple of simple synthetic vector fields. Only windows were updated, while displacements were kept fixed after initialisation. If not mentioned otherwise, in all experiments the initial window radius was 5, the upper radius limit was $r_{\max} = 6$ pixels (px), and no lower limit was imposed. Furthermore we chose $\sigma = 1$ and $e_{\text{outside}} = 0$.

Figure 8 illustrates the adapted windows in the presence of an *affine flow*, i.e. when $\mathbf{u}(x)$ can be written as an affine function in x. If there are no further influences (such as the boundary terms in (b)), the resulting shapes only depend on the Jacobian of $\mathbf{u}$, more precisely on its outer product.

In *transition zones*, e.g. where flows of different direction meet, fixed windows are disadvantageous, because they smooth over high velocity gradients and thus wipe out details. Thus, we investigate this situation in four simplified scenarios. It becomes clear, that the window adaptation is invariant against a constant offset (Fig. 9a vs. 9b) and rotation (Fig. 9a vs. 9c) of the vector field, but only depends on the orientation of the gradient. The sinusoid-shaped vector field in Fig. 9c is motivated by the data set introduced in Fig. 6.

Finally, we investigate scenarios with *sharp motion boundaries* where it is of great importance for the measurement accuracy that the window adaptation process respects the flow discontinuities. The experiments in Fig. 10 combine round and square shaped boundaries with constant and affine vector fields. In any case, the adapted window shapes respect the boundaries well.

C. Synthetic PIV Benchmark Data Set

The data set shown in Fig. 6 and already investigated in Sect. IV-A was created for the *PIV Challenge 2005* and used to evaluate the spatial resolution of 19 PIV algorithms.

First experiments showed that it is essential to have a good initialisation for the proposed adaptive approach. To this

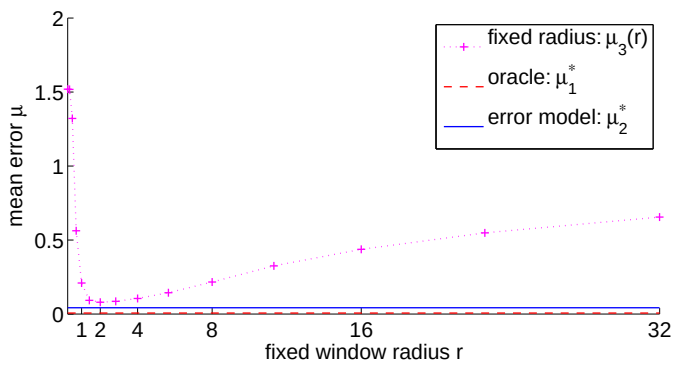

(a) Sinusoid I: fixed radius

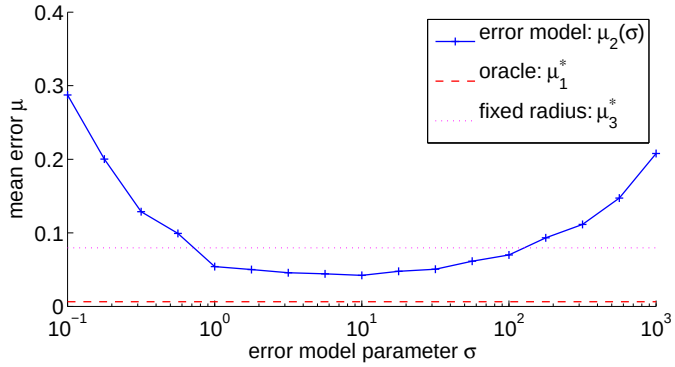

(b) Sinusoid I: error model

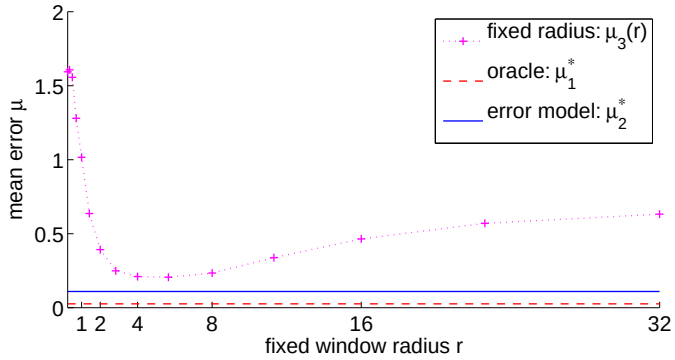

(c) Sinusoid II: fixed radius

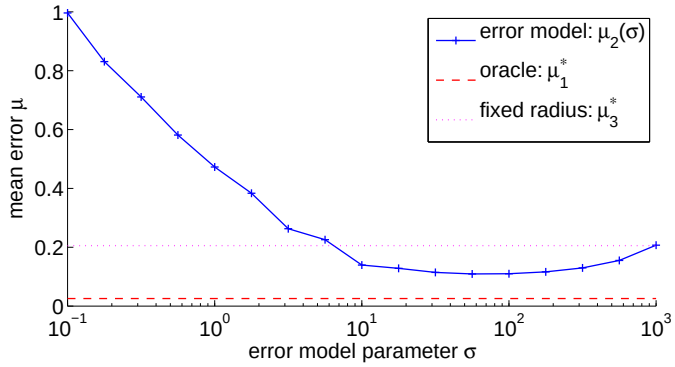

(d) Sinusoid II: error model

Fig. 7. Comparison of window selection strategies: mean error of the fixed radius (dotted line), oracle (dashed line) and the error model based strategy (solid line). **(a)** Results for *Sinusoids I*, comparing the fixed radius $\mu_3(r)$ strategy for varying radii to the best values of the alternative strategies. **(b)** Same as (a), but for the error model strategy and varying parameter σ. Best results are $\mu_1^* = 0.00627$, $\mu_2^* = 0.0421$ (for $\sigma^* = 10$) and $\mu_3^* = 0.0796$ ($r^* = 2$). **(c)-(d):** Same as (a)-(b), but for the data set *Sinusoids II*, with $\mu_1^* = 0.0255$, $\mu_2^* = 0.109$ ($\sigma^* = 10^{7/4}$), $\mu_3^* = 0.206$ ($r^* = 2^{5/2}$).

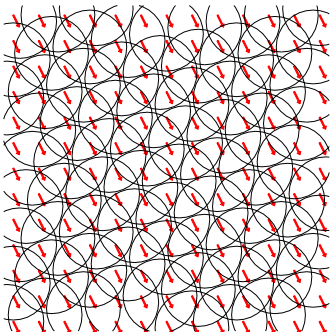

(a) constant vector field

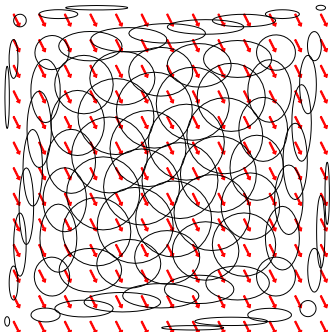

(b) constant vector field, adaptation to image boundaries

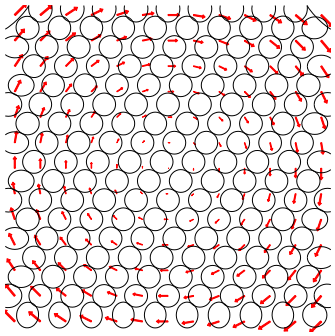

(c) affine vector field with isotropic gradient

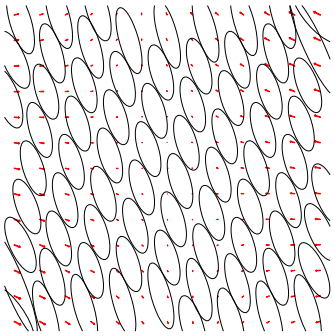

(d) affine vector field with anisotropic gradient

Fig. 8. Synthetic displacement fields (arrows) and some of the adapted windows (ellipses). **Constant vector fields: (a)** The window radii would approach infinity due to the lack of a gradient, but is limited by the constraint $r \leq 6$. **(b)** The windows are additionally constrained to adapt to the boundaries of the image domain by setting $e_{\text{outside}} = 10$. **Affine vector fields: (c)** Rotational field with isotropic gradients leads to round windows. **(d)** Vector field with anisotropic gradients leads to ellipse-shaped windows.

end we first estimated a rough displacement field using fixed windows and then process the data with window adaptation.

For the case *Sinusoids I* the initial vector field was calculated with 5 multiscale levels, using a scaling factor of $\sqrt{2}$ and round windows with radius 6. The adaptive approach used only 3 multiscale levels. We set $\sigma = e_{\text{outside}} = 20$, and constrained the windows radii to the range 2 to 40 px.

The results in Fig. 6 demonstrate that fixed windows can recover the overall structure but smooth over small details. In view of the following processing step, we favour a rough reconstruction over a more detailed but definitely noisy one which can be achieved using smaller windows. With the window adaptation enabled, even the structures at the smallest scale can be reconstructed well up to few outliers, as windows align perpendicular to the velocity gradients and along regions of homogeneous motion. Additional disturbances can

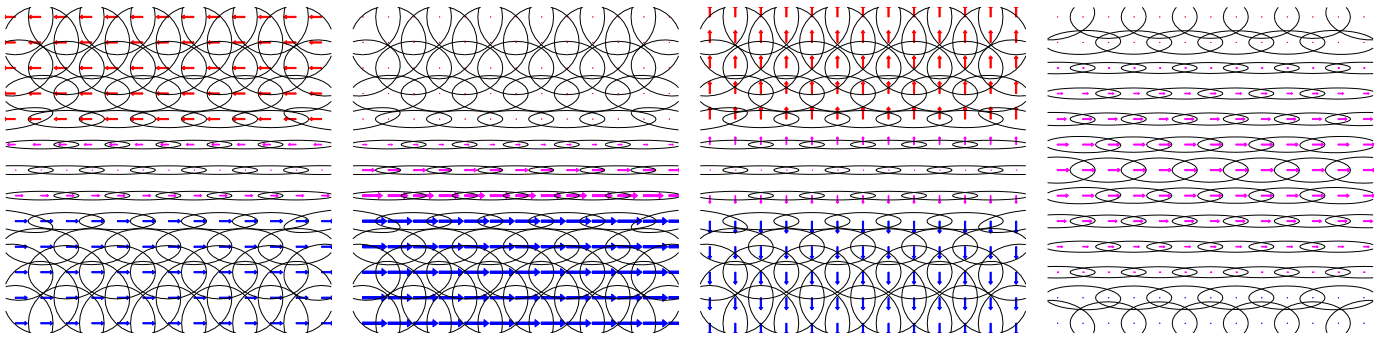

(a) displacement gradient perpendicular to the flow (b) same as (a) with a constant flow superimposed (c) displacement gradient along the flow (d) sinusoid

Fig. 9. Synthetic displacement fields (arrows) and some of the adapted windows (ellipses). Constant vector fields (upper and lower region) enclose a transition zone (middle). **Affine transition zone:** Identical window shapes for different vector fields: **(a)** displacement gradients perpendicular to the flow, **(b)** superimposed by a constant field, and **(c)** gradients parallel to flow. **Sinusoid: (d)** Transition zone is sinus-shaped ($[0, \pi]$). The adaptation scheme aligns the windows perpendicular to the displacement gradient and reduces the size along the transition direction to avoid smoothing out the boundary.

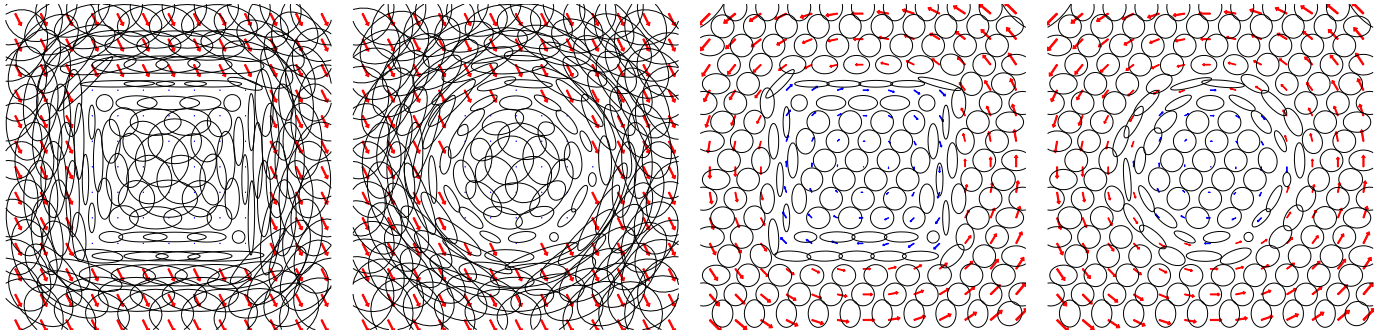

(a) constant vector fields, square boundaries (b) constant vector fields, round boundaries (c) affine vector fields, square boundaries (d) affine vector fields, round boundaries

Fig. 10. Synthetic displacement fields and some of the adapted windows (ellipses), with sharp discontinuities between the inner and outer motion regions. **Constant flows:** Constant flow (red) interrupted by a zero flow (blue) with **(a)** square and **(b)** round inner region. **Rotational (affine) flows:** Two contrarily rotating (affine) flows with **(a)** square and **(b)** round inner region. The adaptation scheme reduces the window sizes near the region boundaries to avoid smoothing over motion discontinuities.

be observed near the upper and lower image boundaries, where windows are extremely compressed.

Additive image noise and unpaired particles in combination with the small structures renders the data set *Sinusoids II* a challenge for any motion measurement algorithm. Again we estimated a coarse description using fixed windows (radius 8 px). Initialised by this result, we run the adaptive approach with the same parameters as for the previous data set, but doubled σ and did not use multiscale calculation.

Figure 12 visualises the resulting displacements and adapted windows. Just as for *Sinusoids I*, the fixed-window approach can only capture the rough motion structures, but the following adaptive approach complements the details even for the smallest wavelength. However, more outliers than for *Sinusoids I* can be observed, where also the adapted windows deviate from their expected vertical alignment.

Finally, we compare our results for *Sinusoids II* to 19 approaches for flow measurement, which were described and benchmarked in [28]. For this purpose we evaluated our method with the same criterion, which is defined as follows: for each sinusoid wavelength λ, we gained a motion profile $u_\lambda(x)$ by averaging the displacements along the vertical axis. Stripes of 10 px at the upper and lower boundary were excluded before. Then using the ground truth profile u^*_λ, the amplitude ratio was calculated as

$$A(\lambda) := \frac{\int_{-\lambda/4}^{+\lambda/4} u_\lambda(x)\,\mathrm{d}x}{\int_{-\lambda/4}^{+\lambda/4} u^*_\lambda(x)\,\mathrm{d}x}\,.$$

The characteristic curve was accurately included into a copy of the comparison plots of the evaluation paper and is presented in Fig. 13. Especially at the lowest wavelength, corresponding to the smallest structures in the data, our adaptive approach outperforms most of the competing implementations.

D. Real Turbulent Experimental Data

Finally, we apply the proposed approach to real PIV data, provided by Johan Carlier in [29] and available at [30]. The experiment describes the turbulent flow behind a cylinder. We chose the image pair number 600 of the data set, recorded with a time difference of 200 μs. Each has a resolution of $1280\,\mathrm{px} \times 1024\,\mathrm{px}$ and dynamics of 12 bit. In Fig. 2a a detailed view of the image data attests the low image quality. An overview over the flow is presented in Fig. 14. As in the

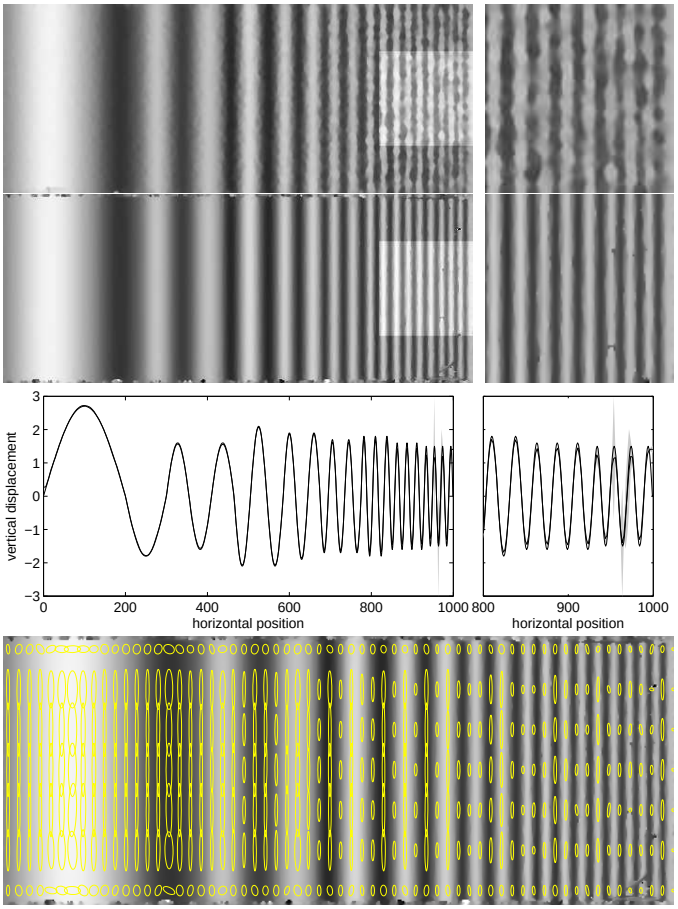

Fig. 11. Synthetic vector field, *PIV-Challenge 2005*, *Case A4*, *Sinusoids I* (no image distortions), full and detailed view of the estimated vertical displacement component (colour map identical to Fig. 6): **Row 1:** Variational correlation with fixed windows ($r = 6$) can only resolve the general structures but smoothes over details. **Row 2:** Using the previous result as initialisation, joint correlation and window adaptation ($\sigma = 20$) can significantly improve the accuracy even for the smallest structures. **Row 3:** statistics (along the complete vertical axis, but 10 px excluded at the upper and lower boundary) of the results in row 2: mean displacement (thick line), range of $\pm$ one standard deviation (shaded gray) and ground truth (thin line, cf. Fig. 6c). **Row 4:** The adapted windows align perpendicular to the velocity gradient.

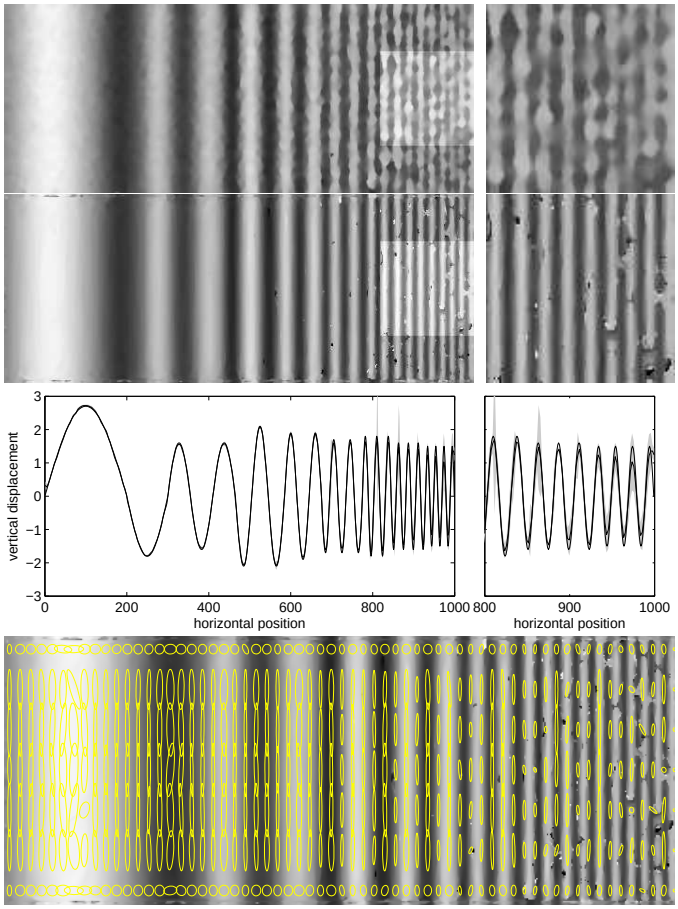

Fig. 12. Synthetic vector field, *PIV-Challenge 2005*, *Case A4*, *Sinusoids II* (with additive image noise and unpaired particles), full and detailed view of the estimated vertical displacement component (colour map identical to Fig. 6): **Row 1:** The high amount of image noise necessitates large windows ($r = 8$) for the approach based on fixed windows. Thus, details are smoothed out, but the general structures can be resolved. **Row 2:** Using the previous result as initialisation, joint correlation and window adaptation ($\sigma = 40$) can significantly improve the accuracy even for the smallest structures, however interrupted by local outliers. **Row 3:** statistics (along the complete vertical axis, but 10 px excluded at the upper and lower boundary) of the results in row 2: mean displacement (thick line), range of $\pm$ one standard deviation (shaded gray) and ground truth (thin line, cf. Fig. 6c), **Row 4:** The windows align perpendicular to the gradient with exception of the vicinity of outliers.

previous experiments, we calculated a coarse vector field using fixed windows ($r = 30$) on the scales $\{1, 2, 4, 8, 16, 32\}$, see Fig. 15a. Initialised by this result, the adaptive approach delivers a more detailed estimation (Fig. 15b). The windows were constrained to $r \in [3, 50]$, and σ was chosen as 100. Only the finest scale was used.

Lacking ground truth data, we employ a vector field calculated by the *Lavision* (http://www.lavision.de/en/) company using their PIV-software *Davis* as reference (Fig. 15c). Our approach smoothes the displacements in regions of homogeneous motion, while the reference solution exhibits some noise. In turbulent regions, however, the windows are adapted such that gradients in the vector field are prevailed.

In order to separate effects of the window adaptation from the influence of initialisation, we rerun the adaptive method with the same parameter but initialised with the reference solution. The result in Fig. 15d shows the same properties as discussed for the one in Fig. 15b. Furthermore, the coarse structure is almost identical in all three solutions, which suggests their correctness, and shows the robustness of our approach with respect to local optima.

Finally, Fig. 14 marks the location of three regions for which we give detailed views of the vector field together with some of the adapted windows. Figure 3 demonstrates how windows sizes reduce in vicinity of a vortex compared to a homogeneous region. The ability to *continuously* control the window orientation is beneficial, for example around the vortex in Fig. 16a. As already demonstrated in Fig. 9, the windows do not necessarily align with the *direction* of the flow, but with its *gradient* as in Fig. 16b.

V. Conclusion

An adaptive approach to measure motion in PIV image data was presented. It is based on the correlation similarity measure, which has proven to be robust also for noisy data.

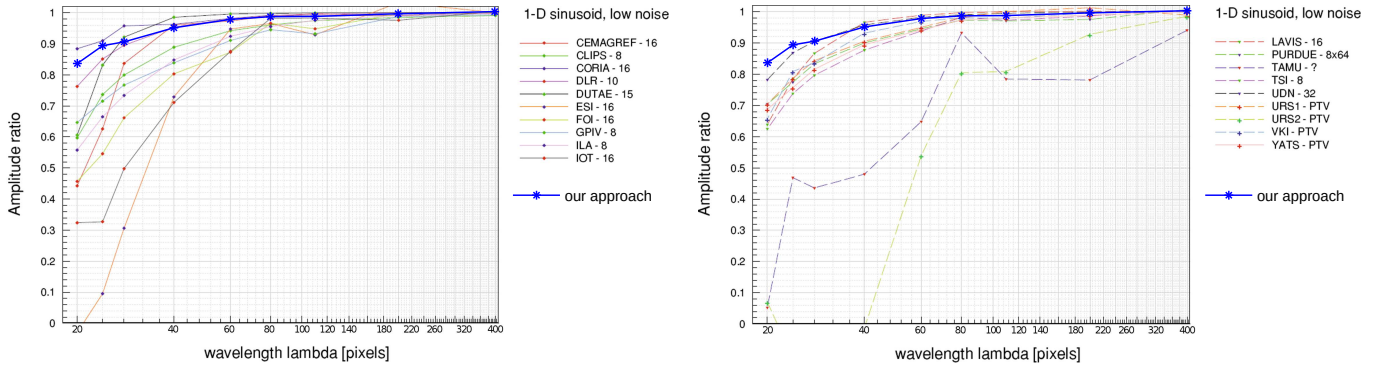

Fig. 13. Synthetic vector field, *PIV-Challenge 2005*, *Case A4*, region *Sinusoids II*: Part 1 and 2 of the comparison of the amplitude response depending on the structure wavelength λ. We included the measurement of our experiments into the plot hard-copied from [28, Fig. 21a] (we redraw the numbers of the horizontal axis for better readability) where 19 flow measurement implementations were compared. Most approaches use cross-correlation while the methods labelled by *CLIPS-8* and *CEMAGREF-16* are based on optical flow. *ESI* and $*$*-PTV* are particle tracking velocimetry methods. For a description of the competing approaches we refer to [28]. Our approach outperforms most of the other implementations, especially the accuracy for very small structures was improved by using adaptive windows.

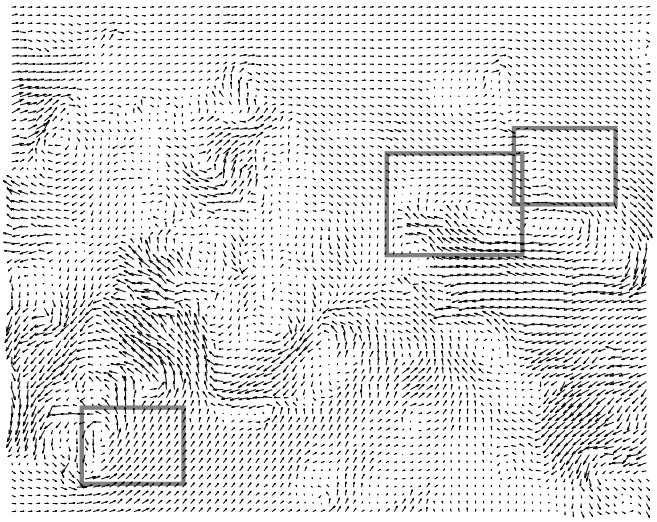

Fig. 14. Real 2D PIV experiment: Overview over the fluid flow, determined by our adaptive correlation method. A mean vector field of about 12 px to the right was subtracted everywhere. The rectangles mark the location of detailed views in – from left to right – Fig. 3, 16a and 16b.

In contrast to classical methods which use a discrete search to find the optimal displacement, we formulate it as a variational problem and use continuous optimisation methods. Furthermore, we employ Gaussian-shaped weighting functions whose shape can be *continuously* controlled and propose a sound adaptation criterion which is based on an error model. Both the displacement measurement and the window adaptation are formulated as two interdependent optimisation problems.

In our experiments we demonstrated the ability of the error model to improve the measurement accuracy and demonstrated the basic behaviour of the adaptation method. We applied our approach to a synthetic PIV benchmark data set and outperformed most of 19 implementations of motion estimators. Finally we showed, that our approach is capable of handling noisy image data from a real experiment. The window adaptation improves the reconstructed vector field in both homogeneous and turbulent regions.

Further work includes to improve the error model function, e.g. to incorporate spatial varying influences (seeding density, image noise level) and further expert knowledge. Regularisation terms could be added to the motion estimation to incorporate prior knowledge, e.g. non-compressibility, on the observed physical process.

Appendix
Derivation of the Noise Term

The term $E_{\text{noise}}(\Sigma)$ in (9) describes the impact of image sensor noise and unpaired particles on the accuracy. For this purpose we assume that the measurement in x is a least-squared solution $\hat{u}$ of independent measurements $\mathbf{u}(y)$, weighted with the same window function as the one used during correlation. For simplicity, and without loss of generality, we assume that the estimation is centred in $x = 0$. Furthermore, here we assume an unbounded variable domain $\Omega = \mathbb{R}^2$.

$$\begin{aligned}\hat{u} :=& \arg\min_{u\in\mathbb{R}^2} \int_{\mathbb{R}^2} w(y,\Sigma)\|u-\mathbf{u}(y)\|_2^2\,\mathrm{d}y \\ =& \frac{\int_{\mathbb{R}^2} w(y,\Sigma)\mathbf{u}(y)\,\mathrm{d}y}{\int_{\mathbb{R}^2} w(y,\Sigma)\,\mathrm{d}y} = \int_{\mathbb{R}^2} \mathrm{G}(\Sigma,y)\,\mathbf{u}(y)\,\mathrm{d}y \qquad (15)\end{aligned}$$

The noise term should only describe the influence of disturbances in the image data, but not the error caused by inhomogeneous motion. Thus, we assume each measurement to be distributed around the true displacement u^*, but disturbed by additive Gaussian noise, i.e. $\mathbf{u}(y) \sim \mathcal{N}\left(u^*, \frac{\sigma^2}{|A|} I\right)$. The constant σ is the relative expected error with respect to the size of the domain A on which a single estimation is based on. Then we define the noise term to be the expected square deviation of (15) from the true solution:

$$E_{\text{noise}}(\Sigma) := \mathcal{E}\left\{\|\hat{u}-u^*\|_2^2\right\} \qquad (16)$$

It is possible to derive a closed form expression for this term. To this end we represent the integral in (15) over an infinite

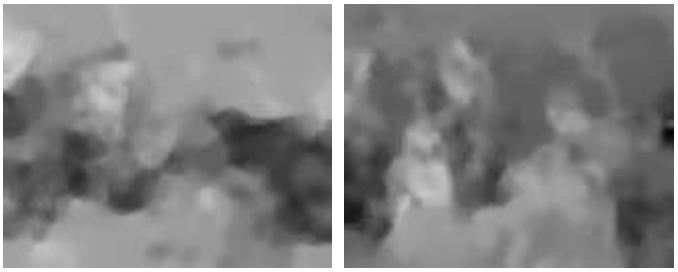

(a) correlation only

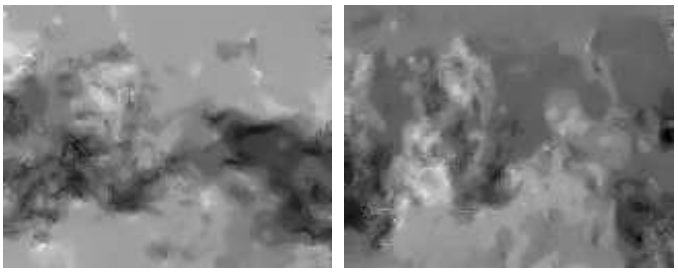

(b) with window adaptation, initialised by (a)

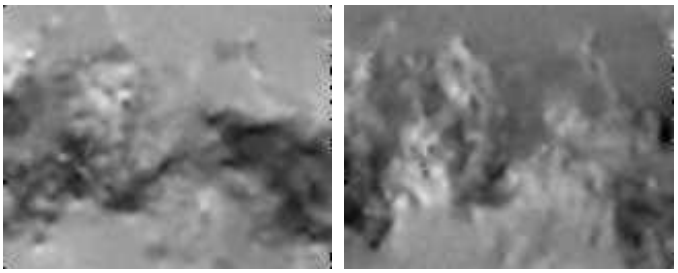

(c) reference solution

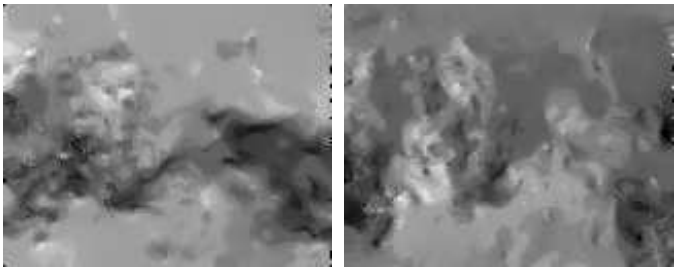

(d) with window adaptation, initialised by (c)

Fig. 15. Real 2D PIV experiment, **Left:** horizontal component of the measured displacements (scale: 8 px (black) to 16 px (white) to the right), **Right:** vertical component (-4 px (black) to +4 px (white) upwards): **(a)** Our approach with fixed windows requires a large support ($r = 30$) to cope with the low image quality and thus smoothes out fine structures. **(b)** Combined correlation and window adaptation ($\sigma = 100$), initialised by (a). In comparison to **(c)** the solution obtained by a commercial correlation software (*Lavision Davis*), our approach preserves details in turbulent regions (middle) and at the same time reduces small disturbances in homogeneous regions (upper and lower boundary) by adapting windows accordingly. **(d)** Initialising our approach with (c) (and using the same parameters as in (b)) shows the robustness with respect to local minima.

domain as a Riemann integral,

$$\hat{u} = \lim_{|A|\to\infty} \lim_{n\to\infty} \hat{u}_n \ ,$$

which is the limit of the Riemann sums $\hat{u}_n$. For simplicity, we assume a square A which is decomposed into $N_n := n^2$ regions of equal size $|A_{ni}| = |A|/N_n$. Sample coordinates are chosen as $y_{ni} \in A_{ni}$. Then we can define

$$\hat{u}_n := \sum_{i=1}^{N_n} |A_{ni}| \mathrm{G}(\Sigma, y_{ni})\, \mathbf{u}(y_{ni}) = \sum_{i=1}^{N_n} w_{ni} \mathbf{u}(y_{ni})$$

with $w_{ni} := |A_{ni}| \mathrm{G}(\Sigma, y_{ni})$. In this formulation, the estimated displacement $\hat{u}_n$ is a linear combination of normally distributed variables and thus is normally distributed as well, i.e. $\hat{u}_n \sim \mathcal{N}(\mu_n, s_n)$ with:

$$\begin{aligned}
\mu_n :=& \mathcal{E}\{\hat{u}_n\} = \sum_{i=1}^{N_n} w_{ni} u^* = u^* \sum_{i=1}^{N_n} |A_{ni}| \mathrm{G}(\Sigma, y_{ni}) \\
s_n :=& \mathcal{E}\big\{(\hat{u}_n - \mu_n)(\hat{u}_n - \mu_n)^\top\big\} = \sum_{i=1}^{N_n} w_{ni}^2 \frac{\sigma^2}{|A_{ni}|} I \\
=& \sigma^2 I \cdot \sum_{i=1}^{N_n} \mathrm{G}(\Sigma, y_{ni})^2 \, |A_{ni}|
\end{aligned}$$

Using $\mathrm{G}(\Sigma, x)^2 = (2\pi\sqrt{\det(2\Sigma)})^{-1} \mathrm{G}\big(\frac{1}{2}\Sigma, x\big)$ (see, e.g., [31, eq. (348)]) we obtain

$$s_n = \frac{\sigma^2}{4\pi\sqrt{\det \Sigma}} I \left(\sum_{i=1}^{N_n} \mathrm{G}\left(\frac{1}{2}\Sigma, y_{ni}\right) |A_{ni}| \right) \ .$$

We assume that for large n the distribution of $\hat{u}_n$ describes the distribution $\hat{u} \sim \mathcal{N}\left(\hat{\mu}, \hat{\Sigma}\right)$ well. Passing the limit, we get the expected result for the mean,

$$\hat{\mu} := \lim_{|A|\to\infty} \lim_{n\to\infty} \mu_n = u^* \lim_{|A|\to\infty} \int_A \mathrm{G}(\Sigma, y)\ \mathrm{d}y = u^* \ ,$$

and – more importantly – the variance

$$\begin{aligned}
\hat{\Sigma} :=& \lim_{|A|\to\infty} \lim_{n\to\infty} s_n = \frac{\sigma^2}{4\pi\sqrt{\det \Sigma}} I \lim_{|A|\to\infty} \int_A \mathrm{G}\left(\frac{1}{2}\Sigma, y\right) \mathrm{d}y \\
=& \frac{\sigma^2}{4\pi\sqrt{\det \Sigma}} I \ .
\end{aligned}$$

Then the definition (16) simplifies to (using [31, eq. (357)])

$$E_{\mathrm{noise}}(\Sigma) = \mathcal{E}\big\{\|\hat{u} - u^*\|_2^2\big\} = \mathrm{tr}\, \hat{\Sigma} = \frac{\sigma^2}{2\pi\sqrt{\det \Sigma}} \ .$$

The noise level σ is the only parameter for this term.

ACKNOWLEDGEMENT

The authors gratefully acknowledge the support provided by the European Commission (Fluid project, FP6-513663) and the German Research Foundation (DFG grant SCHN 457/11-1).

REFERENCES

[1] M. Raffel, C. Willert, and J. Kompenhans, *Particle Image Velocimetry*. Berlin: Springer, 2001.
[2] R. Adrian and J. Westerweel, *Particle Image Velocimetry*, ser. Cambridge Aerospace Series. Cambridge University Press, 2010.
[3] R. J. Adrian, "Twenty years of particle image velocimetry," *Exp Fluids*, vol. 39, pp. 159–169, 2005.
[4] B. Horn and B. Schunck, "Determining Optical Flow," *Artif Intell*, vol. 17, pp. 185–203, 1981.
[5] P. Ruhnau, T. Kohlberger, H. Nobach, and C. Schnörr, "Variational Optical Flow Estimation for Particle Image Velocimetry," *Exp Fluids*, vol. 38, pp. 21–32, 2005.

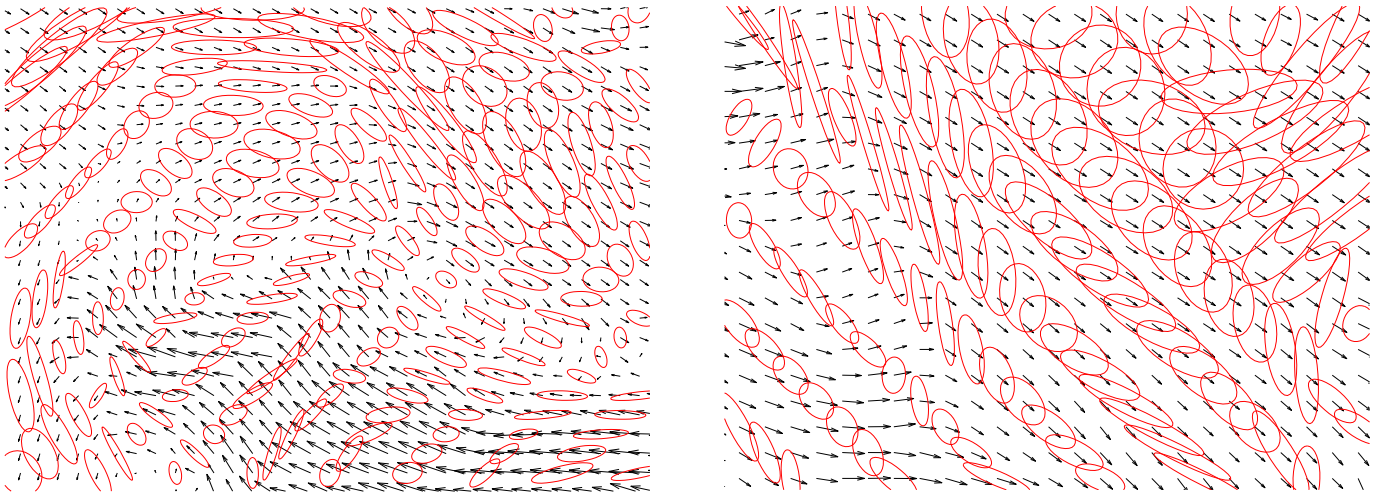

(a) (b)

Fig. 16. Real 2D PIV experiment: Detailed views of Fig. 14, showing the estimated displacements (arrows) and adapted windows (ellipses). **(a)** In the vicinity of a complex flow, the ability to adapt the window orientation *continuously* proves beneficial. **(b)** Windows orientations do not necessarily coincide with the *direction* of the flow, but are perpendicular to the velocity *gradient* here (upper left).

[6] P. Ruhnau, A. Stahl, and C. Schnörr, "Variational Estimation of Experimental Fluid Flows with Physics-Based Spatio-Temporal Regularization," *Meas Sci Technol*, vol. 18, pp. 755–763, 2007.

[7] J. Yuan, C. Schnörr, and E. Mémin, "Discrete Orthogonal Decomposition and Variational Fluid Flow Estimation," *J Math Imaging Vis*, vol. 28, no. 1, pp. 67–80, 2007.

[8] T. Corpetti, D. Heitz, G. Arroyo, E. Mémin, and A. Santa-Cruz, "Fluid Experimental Flow Estimation Based on an Optical-Flow Scheme," *Exp Fluids*, vol. 40, no. 1, pp. 80–97, 2006.

[9] G. Hermosillo, C. Chefd'Hotel, and O. Faugeras, "Variational Methods for Multimodal Image Matching," *Int J Comput Vision*, vol. 50, no. 3, pp. 329–343, 2002.

[10] D. Heitz, P. Heas, E. Mmin, and J. Carlier, "Dynamic Consistent Correlation-Variational Approach for Robust Optical Flow Estimation," *Exp Fluids*, vol. 45, no. 4, pp. 595–608, 2008.

[11] L. Alvarez, C. A. Castaño, M. Garcá, K. Krissian, L. Mazorra, A. Salgado, and J. Sánchez, "3D Motion Estimation Using a Combination of Correlation and Variational Methods for PIV," in *Computer Aided Systems Theory EUROCAST 2007*, vol. 4739/2007, 2007, pp. 612–620.

[12] D. Heitz, E. Mémin, and C. Schnörr, "Variational fluid flow measurements from image sequences: synopsis and perspectives," *Exp Fluids*, vol. 48, no. 3, pp. 369–393, 2010.

[13] J. Nogueira, A. Lecuona, and P. Rodriguez, "Local Field Correction PIV: On the Increase of Accuracy of Digital PIV Systems," *Exp Fluids*, vol. 27, pp. 107–116, 1999.

[14] F. Scarano, "Iterative Image Deformation Methods in PIV," *Meas Sci Technol*, vol. 13, pp. R1–R19, 2002.

[15] J. Nogueira, A. Lecuona, P. Rodriguez, J. Alfaro, and A. Acosta, "Limits on the Resolution of Correlation PIV Iterative Methods. Practical Implementation and Design of Weighting Functions," *Exp Fluids*, vol. 39, pp. 314–321, 2005.

[16] R. Theunissen, F. Scarano, and M. L. Riethmuller, "An Adaptive Sampling and Windowing Interrogation Method in PIV," *Meas Sci Technol*, vol. 18, pp. 265–287, 2007.

[17] ——, "Spatially adaptive PIV interrogation based on data ensemble," *Exp Fluids*, vol. 48, pp. 875–887, 2010.

[18] ——, "On Improvement of PIV image Interrogation near Stationary Interfaces," *Exp Fluids*, vol. 45, pp. 557–572, 2008.

[19] B. Wieneke and K. Pfeiffer, "Adaptive PIV with variable interrogation window size and shape," in *15th International Symposium on Applications of Laser Techniques to Fluid Mechanics*, Lisbon, Portugal, 2010. [Online]. Available: http://ltces.dem.ist.utl.pt/lxlaser/lxlaser2010/upload/1845_qkuqls_1.12.3.Full_1845.pdf

[20] D. Di Florio, F. Di Felice, and G. P. Romano, "Windowing, Re-Shaping and Re-Orientation Interrogation Windows in Particle Image Velocimetry for the Investigation of Shear Flows," *Meas Sci Technol*, vol. 13, pp. 953–962, 2002.

[21] B. D. Lucas and T. Kanade, "An Iterative Image Registration Technique with an Application to Stereo Vision," in *Proceedings of DARPA Imaging Understanding Workshop*, 1981, pp. 121–130.

[22] A. Bruhn, J. Weickert, and C. Schnörr, "Lucas/Kanade Meets Horn/Schunck: Combining Local and Global Optic Flow Methods," *Int J Comp Vis*, vol. 61, no. 3, pp. 211–231, 2005.

[23] F. Becker, "Variational Correlation and Decomposition Methods for Particle Image Velocimetry," Doctoral Thesis, Heidelberg University, Faculty of Mathematics and Computer Sciences, Heidelberg, Germany, August 2009. [Online]. Available: http://www.ub.uni-heidelberg.de/archiv/9766

[24] F. Becker, B. Wieneke, J. Yuan, and C. Schnörr, "A Variational Approach to Adaptive Correlation for Motion Estimation in Particle Image Velocimetry," in *Pattern Recognition*, ser. LNCS, G. Rigoll, Ed., vol. 5096. Springer, 2008, pp. 335–344.

[25] ——, "Variational Correlation Approach to Flow Measurement with Window Adaption," in *14th International Symposium on Applications of Laser Techniques to Fluid Mechanics*, Lisbon, Portugal, 2008. [Online]. Available: http://ltces.dem.ist.utl.pt/lxlaser/lxlaser2008/papers/01.1_3.pdf

[26] R. Theunissen, F. Scarano, and M. L. Riethmuller, "Statistical Adaptivity in PIV Interrogation for Mean Flow Estimation," in *14th International Symposium on Applications of Laser Techniques to Fluid Mechanics*, Lisbon, Portugal, 2008, 14th International Symposium on Applications of Laser Techniques to Fluid Mechanics. [Online]. Available: http://ltces.dem.ist.utl.pt/lxlaser/lxlaser2008/papers/01.1_1.pdf

[27] M. Unser, A. Aldroubi, and M. Eden, "B-Spline Signal Processing: Part II—Efficient Design and Applications," *IEEE T Signal Proces*, vol. 41, no. 2, pp. 834–848, February 1993.

[28] M. Stanislas, K. Okamoto, C. J. Kähler, J. Westerweel, and F. Scarano, "Main Results of the Third International PIV Challenge," *Exp Fluids*, vol. 45, pp. 27–71, Jul. 2008.

[29] J. Carlier and B. Wieneke, "Report 1 on production and diffusion of fluid mechanic images and data," Cemagref, LaVision, Activity Report, European Project FLUID, Nov. 30 2005. [Online]. Available: bscw.irisa.fr/pub/bscw.cgi/d8709/publis/Deliverable_1.2.pdf

[30] "FLUID project homepage: Deliverables." [Online]. Available: http://fluid.irisa.fr/deliverables-eng.htm

[31] K. B. Petersen and M. S. Pedersen, "The Matrix Cookbook," Oct 2008, version 20081110. [Online]. Available: http://www2.imm.dtu.dk/pubdb/p.php?3274