

A variational approach for particle tracking velocimetry

P Ruhnau¹, C Guetter¹, T Putze² and C Schnörr¹

¹ Department of Mathematics and Computer Science, Computer Vision, Graphics, and Pattern Recognition Group, University of Mannheim, 68131 Mannheim, Germany

² Institute of Photogrammetry and Remote Sensing, Dresden University of Technology, 01062 Dresden, Germany

E-mail: ruhnau@uni-mannheim.de, guetter@uni-mannheim.de, torsten.putze@mailbox.tu-dresden.de and schnoerr@uni-mannheim.de

Received 7 March 2005, in final form 27 April 2005

Published 15 June 2005

Online at stacks.iop.org/MST/16/1449

Abstract

We introduce a novel variational approach for evaluating PTV image pairs and sequences in two and three dimensions. We combine a discrete non-differentiable particle matching term with a continuous regularization term. An advanced mathematical method guarantees convergence to a local minimum. The experimental evaluation shows that our variational method competes with three alternative approaches. We outline the potential of our method for further developments.

Keywords: particle tracking velocimetry, particle matching, variational models, finite elements

1. Introduction

The terms *particle image velocimetry* (PIV) and *particle tracking velocimetry* (PTV) denote established classes of image processing methods for extracting the underlying velocity fields in particle images. PIV methods operate on grey-level images, while PTV approaches determine the flow field by tracking individual tracers [17]. PTV methods are capable of yielding a higher resolution than PIV methods, as it is not necessary to average over regions in the image (i.e., interrogation windows). Furthermore, in 3D, PTV can be supported and combined with stereoscopic analysis and 3D reconstruction, leading to high-resolution 3D3C vector fields.

PTV methods are traditionally either based on nearest-neighbour search with geometrical constraints (using four or more consecutive frames) [8, 6], or on binary-image cross correlation (two frames) [24] which computes the cross correlation between regions around particles in the first and in the second frame. More recent approaches include relaxation methods that analyse the probability of particle matching [2, 15] and genetic algorithms that evaluate different pairing schemes based on local morphology conservation or the constraint of vanishing divergence (for incompressible fluids) [22, 4]. Basically, all these methods have two assumptions in common:

- *Small displacements.* While nearest-neighbour search algorithms directly rely on small displacements from one frame of an image sequence to the next (in proportion to the particle density), binary-image correlation methods and relaxation methods both search for possibly corresponding particle images in a certain ‘tracking range’.
- *Smoothness of motion.* Nearest-neighbour search algorithms assume that a particle changes its motion during an image sequence only smoothly. A similar assumption that tacitly underlies binary-image correlation methods is that the particles within a correlation window move with the same speed (if they do not, the correlation peak is less pronounced and the estimates become less reliable). Finally, using relaxation methods, a matching is considered probable if the movement of particles in a certain region can be reduced to a simple translation.

Recently, a novel class of variational approaches to image fluid analysis has been introduced in [20, 9]. These methods were originally developed in the field of computer vision [7, 14] and modified for the purpose of PIV. The basic idea is not to estimate displacement vectors locally and individually, but to estimate vector fields $\vec{u}$ as a whole by minimizing a suitable functional defined over the entire image section. Such functionals typically comprise two terms: a data term measuring how well two images of a sequence match as




Figure 1. Full circles denote particle positions in the first frame, open circles denote positions in the second frame. Left: simple nearest-neighbour search yields mismatches. Right: nearest-neighbour search followed by regularization with smoothness constraint. In the next iteration, T2 will find the correct match.

a function of the vector field $\vec{u}$ to be estimated, and a smoothness term measuring the variation of $\vec{u}$ in terms of first- or second-order partial derivatives. In principle, the smoothness term enforces coherent vector field structures, making the corresponding post-processing steps in connection with traditional local PIV approaches obsolete. Experimental evaluations showed a competitive performance of variational approaches [19]. Furthermore, advanced parallel and fast implementations using off-the-shelf hardware are available [10, 11].

The objective of this paper is to generalize the class of variational approaches to *particle tracking velocimetry*. To this end, we have to replace the *continuous* data term of variational approaches to PIV with a *discrete* non-differentiable particle matching term for PTV. This raises the problem of minimizing such data terms together with a *continuous* regularization term. We accomplish this with an advanced mathematical method which guarantees convergence to a local minimum of such a non-convex variational approach to PTV.

Figure 1 illustrates the basic behaviour of this new type of variational approach to PTV. On the left, figure 1 depicts a common situation where particle matching by nearest-neighbour search fails. The variational PTV approach presented in this paper is able to avoid, and even to revise, such erroneous local decision through the smoothness term (figure 1, right). A key advantage in our opinion is that all ‘rules’ guiding the matching of particles are encoded by the choice of a smoothness term which, in turn, can be related to physical properties of the underlying fluid, such as low divergence [25]. The physical constraints are thus incorporated *directly* (in contrast to, e.g., the indirect incorporation in genetic algorithm approaches). In the following, we will introduce this novel variational approach to PTV (section 2.1) and the corresponding optimization procedure (section 3). The investigation of different smoothness terms in this context is left for future work. Numerical experiments for benchmark image pairs, a comparison with three alternative approaches, as well as results for real-world image sequences will be presented in section 4. We conclude in section 5 by indicating various extensions to this prototypical approach within the variational framework.

2. General problem formulation

2.1. Basic assumptions and constraints

Let S denote the coordinates of the extracted particles in the first image of an image pair, and T denote the coordinates of

the extracted particles in the second image. Then, we define the distance of a specific particle with coordinates S_i to T by

$$d_T(S_i) := d(S_i, T) = \inf_{T_i \in T} d(S_i, T_i),$$

where $d(S_i, T_i)$ is just the Euclidean distance. Therefore, the target velocity field $\vec{u}$ (where $\vec{u}_i$ denotes the displacement of particle S_i from frame 1 to frame 2) minimizes the accumulated distance function

$$D(\vec{u}) = \sum_{i=1}^M d_T(S_i + \vec{u}_i) \quad (1)$$

where $\vec{u} = \vec{u}_1, \vec{u}_2, \dots, \vec{u}_M$, and where M is the number of extracted particles in image 1.

Unfortunately, the minimization of (1) is a highly non-convex problem, as *every* other possible matching minimizes the equation as well. The *local* minimum is just the ‘nearest-neighbour’ solution. We define a convex attraction potential as an increasing continuous function that attracts every particle to its closest neighbour:

$$E_{\text{local}}(\vec{u}) = \sum_{i=1}^M \frac{\alpha}{2} (d_T(S_i + \vec{u}_i))^2. \quad (2)$$

Until now, the particles are only attracted to their nearest neighbours and the minimization of (2) is trivial. This is why we have to make an additional assumption about $\vec{u}$. The prototypical assumption that we want to make use of in this paper is the assumption of *smoothness*. We indicate in section 5 that other assumptions (that include, e.g., physical knowledge) are conceivable.

However, rather than considering vector fields that are close to a constant in a small region (the predominant assumption in PTV) we want to rule out too irregular vector fields by minimizing the magnitudes of the spatial (and, in the case of image sequences, spatiotemporal) gradients of $\vec{u}$:

$$E_{\text{global}}(\vec{u}) = \int_{\Omega} \sum_{j=1}^N |\nabla \vec{u}^j(s)|^2 ds. \quad (3)$$

Please note that $\vec{u} = (u^1, u^2, \dots, u^N)^\top$, where N indicates the dimensionality of the problem (N is usually 2 or 3). The integration variable s is for image pairs in 2D $s = (x, y)^\top$, and in 3D $s = (x, y, z)^\top$, where x , y and z denote the spatial coordinates within the domain Ω . For image sequences, it follows $s = (x, y, t)^\top$ in 2D, and $s = (x, y, z, t)$ in 3D, where x , y and z denote the spatial coordinates, and t is the temporal coordinate.




Figure 2. Full circles denote particle positions in the first frame, open circles denote positions in the second frame. Full rectangles denote the current estimate. One particle has not been detected in the second frame. Minimization of (6) necessarily leads to the wrong match.

Equations (2) and (3) can be combined into the variational framework

$$E(\vec{u}) = E_{\text{local}}(\vec{u}) + \lambda E_{\text{global}}(\vec{u})$$

$$= \underbrace{\sum_{i=1}^M \frac{\alpha}{2} (d_T(S_i + \vec{u}_i))^2}_{\text{data}} + \underbrace{\lambda \int_{\Omega} \sum_{j=1}^N |\nabla \vec{u}^j(s)|^2 ds}_{\text{regularization}}, \quad (4)$$

where E_{local} is called a ‘data term’ which incorporates *local* information, and E_{global} is the *global* regularization term. In this work, the so-called smoothness parameter $\lambda \geq 0$ is considered a user parameter that controls the smoothness of the resulting velocity field. If we choose $\lambda = 0$, no regularization is performed. The reconstructed velocity field is therefore just the ‘nearest-neighbour’ solution, as the *locally* optimal solution for every particle in image 1 is the matching with its nearest neighbour in image 2.

2.2. Outlier treatment

An important problem in the PTV analysis is raised by the fact that usually not all the particles are detected correctly. In 2D it may happen that a particle is visible in the first frame, but moves out of the illuminated plane and is therefore not visible, or beneath the threshold, in the second frame. In 3D, additional problems occur when the 3D reconstruction fails, e.g., due to a very high particle density. Further problems arise from particle images tending to coalesce.

We can distinguish between two error scenarios:

- A particle is extracted from the second image, but not from the first image: in this case, the proposed algorithm can still estimate a reliable velocity field, as it searches matches for all particles in the first frame.
- A particle is visible only in the first frame but not in the second frame: in this error case, the nearest-neighbour search (6) of the proposed algorithm will necessarily find the wrong match in *every* iteration (cf figure 2). Through the smoothness term of (7) this error is propagated to the neighbourhood of the erroneous vector.

The strategy that we merely want to adopt is to eliminate vectors that contribute a high energy to (7). This is achieved

through a threshold: we replace the attraction potential of the data term of (6) by a robust potential—a cut-off potential that cuts off points located beyond an adjustable threshold. These outliers are not considered in the regularization step of the current iteration.

However, the *result* of the regularization step is propagated to the outliers: linear interpolation yields the velocity field also at the location of the outliers, the positions of which are updated, as are the positions of the inliers. The idea is that they may be torn below the threshold in case they were wrongly detected outliers.

In order to improve the performance in image regions with high velocity, we start with a high outlier threshold and then slowly increase this threshold: thus, in the first iterations, particles in fast moving regions will tend to be considered outliers, while particles in slowly moving regions will tend to be considered inliers. The idea is that, in the course of several iterations with an attenuating threshold, more and more particles are considered inliers and the estimated velocity field in the high-velocity regions can converge to the correct flow field. Extensive experiments confirmed this behaviour.

3. Optimization and discretization

Note that the implicit data constraint defined by equation (1) is a non-convex function. Thus, retrieving a local minimum of (4) does not imply having found the global optimum.

We use an auxiliary variable approach that represents a sound mathematical framework and guarantees convergence [3]: in a two-stage iterative algorithm, each iteration comprises a local deformation followed by a global regularization. To justify this approach, we modify the energy $E(\vec{u})$ of (4) by introducing an auxiliary variable $\vec{u}_{\text{aux}}$. The above two steps can then be interpreted as alternate minimizations with respect to each of the two variables, the variable of the initial energy $\vec{u}$ and the auxiliary variable $\vec{u}_{\text{aux}}$.

A general formulation of the energy E_{aux} following [3] and based on formula (4), with the extra auxiliary variable $\vec{u}_{\text{aux}} = \vec{u}_{\text{aux},1}, \vec{u}_{\text{aux},2}, \dots, \vec{u}_{\text{aux},M}$, has the form

$$E_{\text{aux}}(\vec{u}, \vec{u}_{\text{aux}}) = \sum_{i=1}^M \left(\frac{1-\alpha}{2} (d_{S+\vec{u}}(S + \vec{u}_{\text{aux},i}))^2 + \frac{\alpha}{2} (d_T(S + \vec{u}_{\text{aux},i}))^2 \right) + \lambda \int_{\Omega} \sum_{j=1}^N |\nabla \vec{u}^j(s)|^2 ds. \quad (5)$$

The first two terms of equation (5) exhibit the auxiliary variable’s role as an interpolate between $S + \vec{u}$ and T . Globally, we can think of the iterative minimization of E_{aux} as a deformation of the current vector field followed by a regularization. The successive minimization of E_{aux} is equal to subsequent minimization of the following two energies E_I and E_{II} , each with respect to a different variable— E_I with respect to $\vec{u}_{\text{aux}}$, and E_{II} with respect to $\vec{u}$:

local deformation

$$E_I(\vec{u}_{\text{aux}}) = \sum_{i=1}^M \left(\frac{1-\alpha}{2} (d_{S+\vec{u}}(S + \vec{u}_{\text{aux},i}))^2 + \frac{\alpha}{2} (d_T(S + \vec{u}_{\text{aux},i}))^2 \right) \quad (6)$$




Figure 3. Delaunay triangulation of the area covered by particles from an image plane. Line intersections denote extracted particle positions.

global regularization

$$E_{II}(\vec{u}) = \sum_{i=1}^M \left(\frac{1}{2} (d_{S+\vec{u}}(S + \vec{u}_{aux,i}))^2 \right) + \lambda \int_{\Omega} \sum_{j=1}^N |\nabla \vec{u}^j(s)|^2 ds. \quad (7)$$

The two equations can be subsequently iterated in the given order until convergence is obtained. Equations (6) and (7) demonstrate how both minimizations are linked by the term $\sum_{i=1}^M (d_{S+\vec{u}}(S + \vec{u}_{aux,i}))^2$. The minimizing $\vec{u}_{aux}$ of E_I can be interpreted as a trade-off between the closeness to $S + \vec{u}$ and the closeness to T . This gives a good direction of displacement and avoids too large deformations of the auxiliary flow field $\vec{u}_{aux}$. The grid generation is performed using a Delaunay triangulation [1] (cf figure 3). For technical details about grid generation and discretization of equation (7), for 2D and 3D image pairs and image sequences, we refer to [18].

4. Experimental evaluation

In this section, we test the variational PTV approach on synthetic and real 2D and 3D data sets. For the 2D case, we report comparisons of our variational particle tracking approach with three other approaches. Before discussing the results in section 4.5, we first describe the data sets used for the comparison (section 4.1), the preprocessing (i.e., particle extraction and 3D reconstruction, section 4.2), the alternative approaches and their corresponding parameter settings (section 4.3), and the quantitative error measures (section 4.4).

4.1. Data

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

- *Synthetic data.* The Visual Society of Japan (VSJ) has published standard images for particle image velocimetry that are freely available on the internet [16]. For 2D data, we will refer to the test image classified as 301 in the VSJ library. It consists of ten frames taken at intervals of 0.005 s; each frame consists of about 4150 particles. It shows the vertical portion of the impinging jet, with a maximum velocity of 10 pixels/frame. Figure 4 shows the first image from this series along with the correct motion field. We will analyse our 3D approach using the




Figure 4. VSJ standard image 301.




Figure 5. 3D velocity field of VSJ standard image 331.

test images classified as 331 in the VSJ library (jet shear flow). Figure 5 shows a plot of its velocity field.

The advantage of the VSJ images is that the underlying motion fields, as well as the particle coordinates, are available so that the evaluation of different approaches, as well as that of different parameter constellations, is possible. By basing our computations on these particle position data, we have to deal with very high particle concentrations (approximately 4150 particles to be tracked in the 2D case and 3500 particles in the 3D case). We want to evaluate the performance of our algorithm in cases of high particle concentrations, as up-to-date CCD cameras yield increasingly high resolutions, and thus an up-to-date tracking system must be capable of managing high particle concentrations.

In order to achieve a more realistic test scenario, we will randomly delete particles in order to emulate typical individual-particle extraction errors.

- *Real-world data.* Figure 6 shows different frames from a visualized air flow. The camera used is a high-speed camera with 1000 fps. The flow is visualized by styrofoam particles or mirco-balloons. One has to visualize the flow in a way that the mean displacement is approximately 10 pixels/frame.

To analyse the 3D capability of our algorithm, we took the ‘stirred aquarium’ sequence from [12]. It investigates the water flow in a channel made of glass. The velocity of the real flow in the glass channel averaged 30–50 cm s⁻¹. The camera system, which contains three video cameras that operate with 25 fps, was moved continuously in the direction of the flow to optimize the tracking. In order to get the characteristic flow, one has to consider the bias of the moved camera system.




Figure 6. Real-world image: four frames have been superimposed to visualize the overall motion.

For successful processing, a compromise between camera frame rate, exposure time, flow velocity and illumination is necessary. Fast particles detected when using a longer exposure time appear as streaks. This means a loss of accuracy and it generates ambiguities. Longer exposure times also reduce the maximum frame rate. Low frame rates in turn increase errors due to the curvature of the particle paths. High frame rates are correlated with short exposure times, which results in dim particles. Especially, particles at the boundaries of the illumination corridor cannot be segmented by an overall threshold.

In order to get sufficient results in terms of a successful tracking, the highest possible frame rate is required. The more turbulent the flow is, the higher the time resolution has to be in order to get correct matching outputs. The maximum time delay between two epochs depends on the feasibility of the temporal matching. The latter in turn depends on the homogeneity, the turbulence and the velocity of the flow and, of course, on the performance of the matching algorithm and its ability to incorporate spatiotemporal constraints on homogeneity.

4.2. Preprocessing steps

4.2.1. Individual-particle detection. In order to track individual particles, these particles first have to be extracted out of the grey-value structure of the image. Many authors have concentrated on this topic (a comparison of different particle detection approaches can be found in [15]).

While we omitted the particle detection and 3D reconstruction steps in the synthetic cases by directly basing the tracking algorithm on the 3D coordinates provided, we used the so-called particle mask correlation method described in [5, 21] in the 2D real-world cases. For the 3D real-world case, a region growing approach with a discontinuity parameter that divides overlapping particles was used [13].

4.2.2. 3D reconstruction. For 3D PTV, the illuminated scene is usually recorded by a system of three or more cameras.

If the camera calibration is known, then, by use of epipolar geometry, the particle positions can be described by a set of lines. The particles are situated around the intersection points of these lines. According to the accuracy of the calibration and the particle density there will be ambiguities. In some cases it is not possible to resolve them [12]. As the presented algorithm is capable of handling outliers (cf section 2.2), we will, in these cases, consider all *possible* particle locations that cannot be ruled out—anticipating that the wrong candidates will be considered as outliers by the algorithm.

4.3. Approaches and parameter settings

The data sets described above are evaluated with the use of the following approaches and parameter settings.

- **Variational approach (VAR).** The particle coordinates are normalized so that all particles lie between 0 and 1 in all spatial dimensions; the temporal dimension is numbered in integer steps ($t = 1, 2, \dots, T$). For all the test cases we use a smoothness parameter of $\lambda = 0.1$. The parameter α is set to 0.8. In the first iteration, 75% of the particles are considered as outliers and in every iteration 0.1% particles in addition are considered as inliers. No additional particles are considered as inliers if the outlier threshold reaches 0.01. The iteration stops if no further decrease in energy occurs.
- **Four-frame in-line tracking method (FIT) [6, 8].** The movement of the tracer particles is traced frame by frame while checking the geometrical consistency of every possible particle path. Therefore an iterative procedure of firstly the extrapolation of the particle displacement, and secondly the search for the nearest neighbour is implemented. As this method asks for four consecutive frames to track the particles, we will use all four of the frames of the VSJ standard images.
- **Binary-image cross-correlation method (BCC) [24].** This method is considered a variation of the standard cross-correlation PIV, in which the correlation functions are computed for each interrogation window which is centred on the first-frame particles. An adaptive shifting scheme is used.
- **Relaxation method (NRX) [2, 15].** This analysis is based on the probability of particle matching between the first and second frames, defined for every possible pair of particles including the probability of there being no match. A high probability of matching is assumed if the neighbouring particles move similarly.

4.4. Error measures

In this work, we concentrate on two error measures, *yield* and *reliability*:

- **Yield (E_Y)** is the measure of the number of correct vectors produced between two images (n) divided by the total number of particle pairs known to exist between the two images (v):

$$E_Y = \frac{n}{v}. \quad (8)$$




Figure 7. Estimated velocity field VSJ image 301 (left). Two likely error constellations (right): one particle has not been extracted in frame 2, the matching is performed with a close neighbour of this vanished particle (top). Due to three dimensionality of the velocity field, two particles ‘cross’ in the two-dimensional projection. The two-dimensional variational approach presumes the smoothness of the projection and chooses the wrong match (bottom).

- Reliability (E_R) is the measure of the number of correct vectors that were reconstructed by the tracking method (n), divided by the total number of vectors determined by the tracking method (d):

$$E_R = \frac{n}{d}. \quad (9)$$

It is apparent that we can influence both error measures by the chosen parameters: if we use a high outlier threshold, we can expect good reliability (as only matches that fit the model very well are considered), while E_Y will definitely drop. A lower threshold will lead to an increase in E_Y , while decreasing the reliability.

4.5. Numerical results and discussion

4.5.1. 2D results.

Synthetic images. The first test case is the computation of the velocity field between the frames 0 and 1 of the VSJ 301 image sequence. After 700 iterations the solution presented in figure 7 (outlier ratio 3%) is generated.

In the test case, 4042 particles are visible in both images; 4039 matchings are computed; 3894 of which are correct.

This corresponds to a yield rate of $E_Y = 96.34\%$ and a reliability rate of $E_R = 95.93\%$. Figure 8 shows these two error measures through the iteration process. The average angular error of our estimated vectors is 0.24° and the root mean squared (RMS) error is 0.0261 pixel, which would suggest that the performance of our approach is one order of magnitude more exact than that of (cross-correlation-based) PIV techniques. However, these numbers are misleading: when using highly accurate matching techniques, the overall RMS or angular error will be largely caused by inaccuracies in particle extraction.

Figure 10 points up our outlier strategy. As explained in section 2.2, we start using a very low outlier threshold, considering only 25% of the particles as inliers. These particles are located mainly in the left part of the image (top left figure). The velocity is small and the velocity field is smooth, so that these particles fit the model best. In the course of iterations, more and more particles are considered as inliers, so that the reconstructed velocity field resembles the true velocity field even at locations where the velocity is very high. Beginning with iteration 400, more and more particles in the lower right part of the image are considered as inliers and find their correct counterparts (cf figure 10). The velocity induced by these correct matchings, in turn, is propagated through the smoothness constraint to the middle of the image where it raises the number of correct matches. This is why reliability strongly increases between iterations 450 and 500 (cf figure 8).

Figure 9 shows that our approach is rather insensitive to changes of the smoothness parameter λ . However, as λ goes to zero, the reconstructed velocity field approaches the ‘nearest-neighbour solution’ (cf section 2.1) and the accuracy drops. In contrast, if we apply a very large smoothness parameter, the algorithm will not be able to deal with spatial and temporal motion variations, many particles will be considered as outliers and the performance will decrease.

Please note that we used *the same* λ in *all* our experiments—if we had adapted the parameter manually for every experiment we could have achieved better results than those presented in this paper. Changes of the parameter α had only very little influence on the resulting matching.

Table 1 compares the results achieved with our approach, with the results of the approaches introduced in section 4.3.




Figure 8. Graph of the two error measures from iteration 1 to 720.




Figure 9. Graph of the two error measures for different smoothness parameters.




Figure 10. Estimated velocity field after 50 (top left), 300 (top right), 500 (bottom left) and 720 (bottom right) iterations.

Table 1. Comparison of four PTV algorithms: four-frame in-line tracking (FIT), binary-image cross correlation (BCC), relaxation (NRX), variational approach (VAR).

Algorithm	Frames	Estimated matches	Correct matches	Reliability (%)
FIT	0, 1, 2, 3	630	559	88.73
BCC	0 $\rightarrow$ 1	860	788	91.62
	0 $\rightarrow$ 2	863	691	80.07
NRX	0 $\rightarrow$ 1	808	788	97.52
	0 $\rightarrow$ 2	714	680	95.24
VAR	0 $\rightarrow$ 1	872	865	99.20
	0 $\rightarrow$ 2	904	885	97.90

In order to guarantee a fair comparison, we have not used the correct particle coordinates provided by the VSJ, but extracted particle positions by using the particle mask correlation method (cf section 4.2.1). Therefore the number of particles is clearly lower than in the preceding computations.

The variational PTV method finds the largest number of matches while additionally yielding the highest reliability.

4.5.2. 2D plus time results (2D image sequences). The next step is the additional exploitation of temporal smoothness information. Therefore we have to analyse the whole VSJ 301 image sequence consisting of 10 frames. Figure 11 shows the computed trajectories. Table 2 shows the parameters we use and the results that we achieve. Furthermore, the results of the analysis of image pairs only are indicated. In every frame, the computation based on the whole sequence is at least as good as the image pair result. This had to be expected, as additional information is available in the sequence case. The reason why only slight improvements are achieved has already been addressed: we analyse a 2D projection of a 3D velocity field, therefore the smoothness assumption does not necessarily hold at every point in the image. This is why we will later turn to three-dimensional problems.




Figure 11. Computed trajectories from sequence VSJ 301.

Table 2. Error measures for VSJ standard image 301.

Frames	$\alpha = 0.8, \lambda_{sp} = 0.1, \lambda_{tmp} = 10$		$\alpha = 0.8, \lambda = 0.1$	
	Image sequence		Image pairs (%)	
	Yield	Reliability	Yield	Reliability
00 → 01	97.72	96.41	96.34	95.93
01 → 02	97.62	96.61	96.83	95.83
02 → 03	97.69	96.45	97.05	95.81
03 → 04	97.64	96.64	96.90	95.90
04 → 05	97.32	96.63	97.05	96.35
05 → 06	97.64	96.70	97.11	96.18
06 → 07	97.28	96.10	92.99	91.86
07 → 08	97.28	96.30	93.11	92.10
08 → 09	96.33	95.24	93.94	92.87

Real-world images. Figure 12 shows the computed trajectories for the four frames of the 2D real-world image. Visual comparison of the extracted velocity field and image pair suggests that there has not been a wrong match.

4.5.3. 3D results.

Synthetic data. First we compute the 3D velocity field between frames 0 and 1 of the VSJ 331 image sequence. The solution that was generated after 750 iterations is presented in figure 13. In this test case, 3364 particles are visible in both images and 3372 matchings are computed. These matches include all exact matches, and eight particles that do not have a counterpart in the second image but are erroneously matched to another particle. As expected, the 3D results are much better than the 2D results. Computations with volume coordinates of the VSJ 301 sequence show that we achieve matching rates very close to 100% in these test cases, too.

Please note, however, that in real-world scenarios, errors in 3D imaging and 3D reconstruction will lead to missing particles or erroneous particle locations. This is why we want to test the robustness of our algorithm: table 3 shows that even an increase in outlier probability does not deteriorate the results significantly. In these test cases, the indicated percentage of particles (first column) has been randomly removed from *both* images to simulate problems in particle extraction and 3D




Figure 12. Left: computed trajectories from real-world image sequence. Right: velocity vectors between frame 1 and frame 4 with the mean flow component subtracted; amplified for perceivability.




Figure 13. Estimated 3D velocity field for sequence VSJ 331.


reconstruction. The second column indicates the number of particles that are visible in both frames, columns 3 and 4 show the two performance measures.

In order to assess the limits of our approach we consider only every second image. The results indicated in table 3 show that the error measures are still very good. When considering only every third image, however, the approach is no longer able to determine a valid velocity field. In fact, both yield and reliability drop to 0% (i.e., not a single velocity vector is recovered correctly). The algorithm does not find a starting point as the offsets at *every* position in the image are so high that *no* particle is able to find its counterpart in the first iteration and thus the algorithm converges to the wrong minimum. This drawback has to be expected as we are minimizing a highly non-convex functional (cf equation (6)).

Even in this case, perfect matching can be found, if we provide the algorithm with a good initial guess (e.g., by

Table 3. Error measures for VSJ standard image 331.

Removed particles (%)	00 → 01			00 → 02		
	Possible matches	Yield (%)	Reliability (%)	Possible matches	Yield (%)	Reliability (%)
0	3364	100.00	99.76	3192	99.97	99.47
5	3037	100.00	99.84	2881	99.86	99.45
10	2731	100.00	99.60	2586	99.38	99.34
15	2440	100.00	99.59	2307	98.22	99.60
20	2170	100.00	99.40	2053	98.30	99.56
25	1885	100.00	99.74	1809	44.83	44.81
30	1649	100.00	99.40	1557	38.79	39.35
35	1403	99.93	99.64	1339	31.14	31.17
40	1211	100.00	99.26	1131	32.98	33.01

**Figure 14.** Estimated 3D trajectories for the real-world sequence ‘stirred aquarium’. Every third trajectory has been printed.

specifying the overall image velocity, or one single match); we will point out this possibility in section 5.

Real-world data. Figure 14 shows the extracted trajectories from the real-world sequence ‘stirred aquarium’ [12]. The sequence consists of 31 exposures of the whole volume with a three-camera set-up. 1300–1400 particles were detected in every image. 3D reconstruction yielded a total number of 28818 particles from the sequence (i.e., ≈ 930 particles/frame). We used the same parameters used for the synthetic experiments. 22 485 matches were found when using our variational algorithm (i.e., ≈ 750 matches/image pair). The position of each vector is expressed in the initial camera coordinate system.

5. Conclusion and further work

We have introduced and successfully evaluated a variational particle tracking velocimetry approach that combines a discrete matching term and a continuous regularization term. This novel approach can handle 3D image sequences and outperforms standard PTV methods.

A decisive advantage of the variational approach (4) is its potential for further development. Note that the regularization term (second term in (4)) does not incorporate specific physically motivated prior knowledge. As the particle

matching is performed in the 3D object space, this term could be replaced by a term that is better suited for specific tasks: one could, e.g., guarantee zero divergence (in the case of incompressible fluids) and/or the compliance with the Navier–Stokes equations.

Furthermore, our approach could be used for particle tracking in an integrated PIV–PTV method (*super resolution analysis*, cf [23]). One could also add the possibility of specifying border conditions for regions where the correct flow is known.

Acknowledgment

Support by the Deutsche Forschungsgemeinschaft (DFG, SCHN 457/6) within the priority programme ‘Bildgebende Messverfahren in der Strömungsmechanik’ (SPP-1147) is gratefully acknowledged.

References

- [1] Aurenhammer F 1991 Voronoi diagrams—a survey of a fundamental geometric data structure *ACM Comput. Surv.* **23** 345–405
- [2] Baek S J and Lee S J 1996 A new two-frame particle tracking algorithm using match probability *Exp. Fluids* **22** 23–32
- [3] Cohen L D 1996 Auxiliary variables and two-step iterative algorithms in computer vision problems *J. Math. Imaging Vis.* **6** 59–83
- [4] Doh D H, Hwang T G and Saga T 2004 3D-PTV measurements of the wake of a sphere *Meas. Sci. Technol.* **15** 1059–66
- [5] Etoh T and Takehara K 1999 A study on particle identification in particle image velocimetry—particle mask correlation method *J. Vis.* **1** 313–23
- [6] Hassan Y and Canaan R 1991 Full-field bubbly flow velocity measurements using a multiframe particle tracking technique *Exp. Fluids* **12** 49–60
- [7] Horn B and Schunck B 1981 Determining optical flow *Artif. Intell.* **17** 185–203
- [8] Kobayashi T, Saga T and Segawa S 1989 Multipoint velocity measurement for unsteady flow field by digital image processing *J. Vis.* **5** 197–202
- [9] Kohlberger T, Mémin E and Schnörr C 2003 Variational dense motion estimation using the Helmholtz decomposition *Scale Space Methods in Computer Vision (Lecture Notes in Computer Science vol 2695)* ed L D Griffin and M Lillholm (Berlin: Springer) pp 432–48
- [10] Kohlberger T, Schnörr C, Bruhn A and Weickert J 2003 Domain decomposition for parallel variational optical flow computation *Pattern Recognition, Proc. 25th DAGM Symposium vol 2781* ed B Michaelis and G Krell (Berlin: Springer) pp 196–203

- [11] Kohlberger T, Schnörr C, Bruhn A and Weickert J 2004 Parallel variational motion estimation by domain decomposition and cluster computing *Proc. ECCV 2004* vol 3024 ed T Pajdla and J Matas (Berlin: Springer) pp 205–16
- [12] Maas H-G 1992 Complexity analysis for the establishment of image correspondences of dense spatial target fields *Int. Arch. Photogramm. Remote Sens.* **29** 102–7
- [13] Maas H-G 1992 Digitale photogrammetrie in der dreidimensionalen strömungsmesstechnik *Dissertation* ETH Zürich, Institut für Geodäsie und Photogrammetrie
- [14] Nagel H H and Enkelmann W 1986 An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences *IEEE Trans. Pattern Anal. Mach. Intell.* **8** 565–93
- [15] Ohmi K and Li H-Y 2000 Particle-tracking velocimetry with new algorithms *Meas. Sci. Technol.* **11** 603–16
- [16] Okamoto K, Nishio S and Kobayashi T 2000 Standard images for particle-image velocimetry *Meas. Sci. Technol.* **11** 685–91
- [17] Raffel M, Willert C E and Kompenhans J 1999 *Particle Image Velocimetry. A Practical Guide* (Berlin: Springer)
- [18] Ruhnau P, Gütter C and Schnörr C 2004 A variational approach for particle tracking velocimetry *Comp. Science Series, Technical Report* Dept. Math. and Comp. Science, University of Mannheim, Germany
- [19] Ruhnau P, Kohlberger T, Nobach H and Schnörr C 2004 Variational optical flow estimation for particle image velocimetry *Lasermethoden in der Strömungsmesstechnik* vol 12 (Karlsruhe: Deutsche Gesellschaft für Laser-Anemometrie GALA e.V.) pp 30/1–30/8
- [20] Ruhnau P, Kohlberger T, Nobach H and Schnörr C 2005 Variational optical flow estimation for particle image velocimetry *Exp. Fluids* **38** 21–32
- [21] Saga T, Segawa S, Kobayashi T and Hu H 2000 Development and evaluation of an improved correlation based PTV method *Proc. 6th Triennial Symp. on Fluid Control, Measurement and Visualization (Sherbrooke, Canada)*
- [22] Sheng J and Meng H 1998 A genetic algorithm particle pairing technique for 3d velocity field extraction in holographic particle image velocimetry *Exp. Fluids* **25** 461–73
- [23] Stitou A and Riethmuller M L 2001 Extension of PIV to super resolution using PTV *Meas. Sci. Technol.* **12** 1398–403
- [24] Uemura T, Yamamoto F and Ohmi K 1989 A high speed algorithm of image analysis for real time measurement of two-dimensional velocity distribution *Flow Visualization* ed B Khalighi, M Braun and C Freitas (New York: ASME) pp 129–33
- [25] Yuan J, Ruhnau P, Mémin E and Schnörr C 2005 Discrete orthogonal decomposition and variational fluid flow estimation *5th Int. Conf. on Scale Space and PDE Methods in Computer Vision* pp 267–78