

Volumetric Flow Estimation for Incompressible Fluids using the Stationary Stokes Equations

Katrin Lasinger¹

¹ ETH Zurich

Christoph Vogel²

² Graz University of Technology

Konrad Schindler¹

Abstract

In experimental fluid dynamics, the flow in a volume of fluid is observed by injecting high-contrast tracer particles and tracking them in multi-view video. Fluid dynamics researchers have developed variants of space-carving to reconstruct the 3D particle distribution at a given time-step, and then use relatively simple local matching to recover the motion over time. On the contrary, estimating the optical flow between two consecutive images is a long-standing standard problem in computer vision, but only little work exists about volumetric 3D flow. Here, we propose a variational method for 3D fluid flow estimation from multi-view data. We start from a 3D version of the standard variational flow model, and investigate different regularization schemes that ensure divergence-free flow fields, to account for the physics of incompressible fluids. Moreover, we propose a semi-dense formulation, to cope with the computational demands of large volumetric datasets. Flow is estimated and regularized at a lower spatial resolution, while the data term is evaluated at full resolution to preserve the discriminative power and geometric precision of the local particle distribution. Extensive experiments reveal that a simple sum of squared differences (SSD) is the most suitable data term for our application. For regularization, an energy whose Euler-Lagrange equations correspond to the stationary Stokes equations leads to the best results. This strictly enforces a divergence-free flow and additionally penalizes the squared gradient of the flow.

1. Introduction

The basis of experimental fluid dynamics is to observe the flow in a volume of fluid. In order to densely visualize the flow of the (transparent) fluid in the whole volume, one can use a method called *Particle Imaging Velocimetry* (PIV) [1, 31]: Tracer particles are injected into the fluid and their motion is recorded with multiple high-speed cameras. The 3D setup is illustrated in Fig. 1 and an exemplary video frame is shown in Fig. 2.

Applications include the analysis and design of hydro-




Figure 1. Volumetric flow estimation. From 3D particle distributions at consecutive time-steps (*top*), reconstruct a dense motion field \mathbf{V} . xyz -components of the estimated flow field (*bottom*).

dynamically efficient vehicles and machines, or behavioral studies of aquatic organisms that live in flowing water. While early variants of PIV operated in a 2D plane (implemented by illuminating only a thin slice of the volume with a laser light-sheet), the trend is to observe fluid flow in 3D. The (synchronized) multi-view video is processed in two steps. First, per time-step, all views are used to recover the 3D particle distribution, often using variants of space-carving (so-called *tomographic particle imaging velocimetry*, or Tomo-PIV). Second, motion vectors are estimated by dense matching between consecutive 3D reconstructions. From the perspective of computer vision, the matching step employs rather simplistic methods, essentially exhaustive winner-takes-all matching of local 3D windows, followed by some sort of heuristic smoothing. Here, we ask the obvious question: *Can we improve PIV with modern optic flow estimators, while still handling realistic, high-resolution fluid dynamics data?*

We base our approach on a 3D extension of the primal-dual solution [5, 28] of variational flow estimation [17]. Modeling the problem in the 3D domain allows us to exploit the physical properties of our scenario for regularization, which no longer hold in a 2D image projection. Here, we prefer a variational approach and first-order methods [5], because these methods have a manageable memory foot-




Figure 2. Particle image from one view, with zoomed detail.

print and offer the possibility to include the aforementioned physical constraints. As in the 2D case, [7, 36, 37], we prefer to use window-based data terms rather than per-pixel brightness constancy. For our (soft, probabilistic) particle volumes, where all particles look the same, we posit that matching information can still be found in the local patterns of the particles.

Standard regularizers, like a quadratic penalty on the gradients of the flow field or total variation (TV), do not consider the special properties of liquids, which are incompressible and thus should have a divergence-free flow field. Regularizers which penalize divergence have been tried for both 2D and 3D problems [2, 13, 14]. Here, we propose an energy whose optimality conditions correspond to the stationary Stokes equations. This energy formulation integrates nicely into our optimization framework and leads to the desired incompressibility constraint on the flow field.

A challenge of 3D flow estimation is the high computational cost, as the number of voxels, respectively variables, increases cubically with the resolution. On one hand, one could argue that lower-resolution reconstruction is sufficient, because the effective resolution of PIV is limited by the particles’ size and density. On the other hand, naively down-sampling the volume will drastically smooth out the local particle likelihood, and thereby destroy both robustness against mismatches and localization accuracy. We therefore propose a semi-dense approach, where flow vectors are computed (and regularized) on a lower-resolution grid, while the data-term is evaluated at full resolution to retain the precision and robustness of the original particle distribution. We demonstrate that the results of the semi-dense approach are comparable to per-voxel flow estimation, but able to handle large volumetric datasets recorded with modern high-resolution sensors. In our experiments, the proposed method is superior to simple local flow estimation and, without any post-processing, delivers results comparable with the best available 3D fluid flow methods.

2. Related Work

In experimental fluid mechanics, *particle imaging velocimetry* (PIV) is the estimation of a dense velocity field in a fluid volume with the help of tracer particles and optical images [1, 31]. Though traditionally focusing on 2D

motion on a single laser sheet, recent PIV approaches also tackle 3D velocity fields [11, 33]. Elsinga *et al.* [11] were the first to handle high particle densities in 3D with their tomographic particle image velocimetry (Tomo-PIV) method. Per time-step, the particle distribution in the volume of interest is reconstructed from multi-view imagery. The velocity field is estimated by subsequent 3D cross-correlation of large local 3D windows (in PIV terminology “interrogation volumes”) containing the reconstructed particles. The large size of the interrogation volumes, determined by the practically viable particle density, limits the spatial resolution of the reconstructed flow field. More recent methods have focused on post-processing multiple consecutive two-view velocity fields with the help of a dynamic model to increase reconstruction quality [25]. Technical approaches include iterative volume deformation with adaptive window sizes [8], methods in the spirit of Lucas-Kanade tracking [6] and reconstructing trajectories of particle patterns over time [21]. Schanz *et al.* [33] propose tracking of individual particles over long time sequences instead. While tracking individual particles is certainly an option to refine PIV results, the basic two-frame case is arguably better captured by dense flow estimation, which more naturally matches the continuous and physically constrained nature of flowing fluids. Furthermore, individual particle tracking recovers the flow only at particle locations. Hence, additional post-processing is needed to interpolate sparse tracks to a grid and to apply the relevant physical constraints [12, 34]. Heitz *et al.* [14] point out potential research directions for variational PIV and advocate the use of correct physical constraints, focusing mainly on 2D problems. Alvarez *et al.* [2] were, to our knowledge, the first to present a variational model of 3D-PIV that accounts for the physical properties of incompressible fluids. However, their method is applicable only as refinement after an initial flow estimation with another (not physically grounded) model. Even so, results are only shown for small toy volumes (up to $256 \times 128 \times 144$ voxels).

Also related is work by Gregson *et al.* [13], who propose a flow reconstruction algorithm for dye-injected two-media fluids, primarily aiming for visually pleasing results over small volumes ($\approx 100^3$ voxels). Their data contains more structural information than our particle images, hence the step from 2D to 3D is less ambiguous. On the other hand, spatial information is available only in areas where dye is visible, thus multiple time steps are needed to densely cover the domain. Technically, they show that the pressure-projection step, commonly used for fluid simulations, is equivalent to a projection on the divergence-free subspace, which they formulate as a proximal operator. Like theirs, our scheme is also motivated by physical properties of fluids, additionally we show that our regularizer *emerges from a proper energy formulation*, naturally derived from the stationary Stokes equations. It is therefore not necessary,

to perform the – computationally expensive – reprojection onto the subspace of divergence-free motion fields in each iteration of the algorithm. Instead, we can include incompressibility as a hard constraint. With modern optimization techniques, this leads to a much more efficient algorithm.

In the field of medical image processing, volumetric flow estimation is used to register 3D scans from different imaging modalities, *e.g.*, computer tomography (CT) and magnetic resonance imaging (MRI). In medical imaging, 3D flow estimation [16, 23, 26, 29] is further important for applications that aim to compensate 3D organ motion in time-resolved medical image sequences (*e.g.*, due to respiration).

Recently, 2D optical flow benchmarks have been dominated by label-based methods [7, 24], propagation methods [4, 18], neural regression networks [10] and models that exploit scene-specific properties like semantics [35, 3]. Most of these models do not scale well to the volumetric domain and struggle heavily with memory consumption. Further, it is not obvious how to incorporate physical properties of the problem into any of them. Finally, our particle data does not provide any semantic, structural or textural cues. Consequently, we feel that a variational approach is the most suitable one for our application scenario.

3. Approach

Our 3D flow estimation pipeline from multiple perspective images is depicted in Figure 3. Starting from ≥ 3 synchronous images of the volume, seeded with tracer particles, we perform a 3D reconstruction with our own implementation of the tomographic reconstruction technique – essentially a soft version of space-carving, which outputs a 3D voxel space of “particle presence probabilities”, see Sec. 3.1. Two such 3D volumes from consecutive time-steps then serve as input for 3D flow estimation. We prefer to work with floating-point scores that indicate the probability of a voxel being occupied, rather than take hard decisions about the presence or absence of particles. Our main contribution, the volumetric flow estimation, is described in Section 3.2. We review the primal-dual approach, introduce data terms, regularizers and complete cost functions for the 3D fluid case, and explain the semi-dense flow estimation.




Figure 3. Flow estimation pipeline. (left) 2D input data, (middle) 3D reconstruction, (right) 3D flow estimation from 2 time steps.

3.1. Tomographic Reconstruction

For 3D reconstruction we follow the Tomo-PIV approach and implement the MART algorithm [15]. Addi-

tionally, after each MART iteration we apply anisotropic Gaussian smoothing ($3 \times 3 \times 1$ voxels) to the reconstructed volume, to account for elongated particle reconstructions along the z -axis due to the camera setup [9]. MART is an iterative solver for the inverse problem

$$\sum_{j \in \mathcal{R}_i} \omega_{i,j} E(X_j, Y_j, Z_j) = I(x_i, y_i), \quad (1)$$

where $I(x_i, y_i)$ denotes the observed pixel intensities, $E(X_j, Y_j, Z_j)$ the unknown voxel intensities and \mathcal{R}_i is a list of all voxels traversed by the viewing ray through pixel (x_i, y_i) . The weight $\omega_{i,j} \in [0, 1]$ depends on the distance between the voxel center and the line of sight and is essentially equivalent to an uncertainty cone around the viewing ray, to account for aliasing of discrete voxels. For each pixel i in every camera, and for each voxel j , the following update step is performed:

$$E(X_j, Y_j, Z_j)^{k+1} = E(X_j, Y_j, Z_j)^k \cdot \left(I(x_i, y_i) / \sum_{j \in \mathcal{R}_i} \omega_{i,j} E(X_j, Y_j, Z_j)^k \right)^{\omega_{i,j}}. \quad (2)$$

The result of the 3D reconstruction is a $N \times M \times L$ voxel space with a scalar intensity value per voxel, where high intensity indicates high likelihood that the voxel is occupied by a particle. With increasing number/density of particles the ambiguity increases, since multiple 3D configurations are plausible that would reproject to the given images. These ambiguities lead to so-called “ghost particles” [22], which cannot be ruled out with the available evidence. Therefore, our flow estimation algorithm will have to deal with particles that do not actually exist. Ghost particles are on average lower in intensity than true particles. However, the imaging setup also leads to intensity variations between true particles, such that ghost particles cannot be filtered by thresholding. We prefer not to heuristically resolve ambiguities at an early stage, and instead use the soft, noisy occupancy probabilities as input for 3D flow estimation.

MART tends to decrease also the probabilities of true particles, see Figure 4. We found it advantageous to apply non-linear contrast stretching to the volume before the subsequent flow computation: $E_{\text{out}} = (E_{\text{in}})^\gamma$, with $\gamma = 0.7$.

3.2. Volumetric Flow

Given two 3D volumes $V_0, V_1: \Omega \rightarrow \mathbb{R}^+$, defined over the volumetric domain $\Omega \subset \mathbb{R}^3$, we aim to reconstruct the 3D motion field $\mathbf{v}: \Omega \rightarrow \mathbb{R}^3$. The functional $\mathbf{v} = (u, v, w)^T$ represents a mapping of points \mathbf{p} in V_0 to points $(\mathbf{p} + \mathbf{v})$ in V_1 .

Like 2D optical flow, the problem is ill-posed when relying only on the local similarity $E_D(V_0, V_1, \mathbf{v})$. The solution is to incorporate prior information into the model, in the form of a regularizer $E_S(\mathbf{v})$. The optimal flow field can




Figure 4. Histogram of true and ghost particle intensities after non-linear contrast stretching.

then be found by minimizing the energy

$$\lambda E_D(V_0, V_1, \mathbf{v}) + E_S(\mathbf{v}) \rightarrow \min_{\mathbf{v}}. \quad (3)$$

To minimize (3), we write the problem as a saddle-point problem and apply the primal-dual algorithm, following [5]. Despite the high frame-rate of the cameras, particles travel multiple voxels between successive frames. To handle these larger motions, we embed the problem in a coarse-to-fine scheme and repeatedly linearize V_1 at the intermediate solution, to obtain a convex approximation of our data term. The coarse-to-fine scheme implicitly uses a larger neighborhood for the data term at coarser pyramid levels.

Historically, optical flow estimation tries to avoid local assumptions about the motion and use pixelwise brightness constancy (BCA) as data cost E_D [17]. However, BCA turns out to not work well for noisy data, and recent work prefers patch-based data terms from the stereo matching literature, *e.g.* [7, 36, 37]. The weakness of BCA is aggravated by the low density of the evidence in PIV data (in our case ≈ 0.0003 particles per voxel). To alleviate this problem, we tried different patch-based data costs and found the sum of squared differences (SSD) in a neighborhood window to work well. We also tried more robust data terms like Census and CSAD [37], but found them to perform worse. It appears that for particle flow setups, the focus lies on high accuracy of the flow vectors, whereas robustness is less crucial due to controlled lighting and high contrast particles. Thus, there is no need to sacrifice sensitivity for robustness of the similarity measure.

We further experimented with different regularizers E_S . In the flow literature, regularization usually amounts to smoothing the flow field. The ancestral smoothness term is quadratic regularization (QR) [17], later it emerged that the more robust Total Variation (TV) [32, 39] was often preferable. Formally, these regularizers are defined as

$$\begin{aligned} \text{TV}(\mathbf{v}) &= \int_{\Omega} |\nabla u| + |\nabla v| + |\nabla w| dx, \\ \text{QR}(\mathbf{v}) &= \frac{1}{2} \int_{\Omega} |\nabla u|^2 + |\nabla v|^2 + |\nabla w|^2 dx. \end{aligned} \quad (4)$$

In our work we choose $|\cdot| := \|\cdot\|_2$. Our main focus lies on

the estimation of flow in fluids, especially water. The incompressibility of water – and many other fluids – implies a divergence-free flow field, it is thus natural to regularize by prohibiting or penalizing divergence. To that end, we investigate the stationary Stokes equations:

$$-\mu \Delta \mathbf{v} + \nabla p = \mathbf{f} \quad \text{subject to } \nabla \cdot \mathbf{v} = 0. \quad (5)$$

Here, the Laplace operator Δ is applied component-wise on the flow field \mathbf{v} . Eq. (5) can be interpreted as follows: An external force field \mathbf{f} acts on the fluid and leads to a deformation. Both, incompressibility and viscosity (viscosity constant μ) prevent the fluid from simply following \mathbf{f} . The pressure field p has to compensate for differences in force and motion field, leading to an equilibrium. It should be mentioned that (5) lacks the transport equations (inertia part) of the full Navier-Stokes model of fluid dynamics. However, the basic optical flow model only looks at 2 frames, and for very short time intervals (5) can serve as reasonable approximation.

The stationary Stokes equations (5) correspond to the Euler-Lagrange equations of the following energy:

$$\min_{\mathbf{v}} \max_p \int_{\Omega} \frac{\mu}{2} (|\nabla u|^2 + |\nabla v|^2 + |\nabla w|^2) + \langle \nabla \cdot \mathbf{v}, p \rangle - \langle \mathbf{v}, \mathbf{f} \rangle dx \quad (6)$$

In this saddle-point problem, the pressure p takes the role of a Lagrange multiplier for the incompressibility condition. In our reconstruction task, the role of the term $\langle \mathbf{v}, \mathbf{f} \rangle$ for the force field is filled by the data term, *i.e.* we can interpret the remaining terms as a regularizer. Eq. (6) suggests to employ quadratic regularization on the flow field, and to apply the incompressibility condition as a hard constraint:

$$\text{QRD}_{\infty}(\mathbf{v}) = \frac{1}{2} \int_{\Omega} |\nabla u|^2 + |\nabla v|^2 + |\nabla w|^2 + \delta_{\{0\}}(\nabla \cdot \mathbf{v}) dx, \quad (7)$$

with δ_C the indicator function of the convex set C . Indeed, this physically motivated regularization scheme nicely fits into our optimization framework, and leads to the best results for our data.

A price to pay is that the regularizer is no longer strongly convex, so we cannot accelerate the optimization as in [5]. We therefore also test a version where the incompressibility constraint is replaced by a soft penalty, while keeping the the quadratic regularization of the gradients:

$$\text{QRD}_{\alpha}(\mathbf{v}) = \frac{1}{2} \int_{\Omega} |\nabla u|^2 + |\nabla v|^2 + |\nabla w|^2 + \alpha (\nabla \cdot \mathbf{v})^2 dx. \quad (8)$$

Estimation Algorithm. We minimize our proposed energy functional (3) in a discrete setting and partition the domain Ω into a regular voxel grid. The objective is to assign a displacement vector $\mathbf{v}_p := \mathbf{v}(p) \in \mathbb{R}^3$ to each voxel $p \in \mathbf{p} = \{1 \dots N\} \times \{1 \dots M\} \times \{1 \dots L\}$, so $\mathbf{v} \in V := \mathbb{R}^{3NML}$. We

briefly review the primal-dual approach [5] for problems of the form

$$\min_{\mathbf{v}} F(D\mathbf{v}) + G(\mathbf{v}), \quad (9)$$

where $D:V \rightarrow Y$ is a linear operator that depends on the form of the regularizer. G denotes the discretized version of the data term from (3) and F one of the investigated regularizers from (4, 7, 8). In case of QR or TV, the linear mapping D implements the finite differences to approximate the spatial gradient of the flow in each coordinate direction, which leads to $Y := Y^1 \times Y^2 \times Y^3 = \mathbb{R}^{3NML} \times \mathbb{R}^{3NML} \times \mathbb{R}^{3NML}$. In order to base our regularizer also on the divergence of the flow field, either as hard (7) or as soft (8) constraint, we extend the linear mapping D accordingly. Through a linear approximation of the 3D divergence based on finite (backward) differences, we arrive at $Y := Y^1 \times Y^2 \times Y^3 \times Y^4$. Here $Y^4 \in \mathbb{R}^{NML}$ holds the dual variables for the incompressibility constraint, respectively penalty; the pressure field (5). For convex F and G we get the primal-dual form

$$\min_{\mathbf{v} \in V} \max_{\mathbf{y} \in Y} \langle D\mathbf{v}, \mathbf{y} \rangle - F^*(\mathbf{y}) + G(\mathbf{v}), \quad (10)$$

where $F^*(\mathbf{y}) := \max_{\mathbf{v} \in V} \mathbf{v}^\top \mathbf{y} - F(\mathbf{v})$ denotes the conjugate of F . In this form the problem can be solved by iteratively updating \mathbf{v} and \mathbf{y} according to

$$\begin{aligned} \mathbf{v}^{k+1} &= (I + \tau \partial G)^{-1} (\mathbf{v}^k - \tau D^\top \mathbf{y}^k) \\ \mathbf{y}^{k+1} &= (I + \sigma \partial F^*)^{-1} (\mathbf{y}^k + \sigma D(2\mathbf{v}^{k+1} - \mathbf{v}^k)). \end{aligned} \quad (11)$$

The data and smoothness terms are decoupled and the updates of the primal and dual variables can be solved pointwise, *c.f.* [5]. In particular, the proximal operator for F^* for TV is given by the pixel-wise projection $\forall p \in \mathbf{p}$:

$$\left((I + \sigma \partial F^*)^{-1} (\mathbf{y}^i) \right)_p = \frac{\mathbf{y}_p^i}{\max(1, \|\mathbf{y}_p^i\|_2)} \quad (12)$$

onto the unit ball for $i=1,2,3$. Similarly, for QR and QRD_α the proximal operator is defined as

$$\left((I + \sigma \partial F^*)^{-1} (\mathbf{y}^i) \right)_p = \mathbf{y}_p^i \frac{\alpha^i}{\alpha^i + \sigma}, \quad (13)$$

with $\alpha^i=1$. In case of QRD_α we additionally have $\alpha^4 := \alpha$. In case we utilize hard incompressibility constraints, elements from Y^4 remain unchanged and are only affected by the explicit gradient steps. For further details refer to [5, 37].

The SSD data cost has the form

$$\text{SSD}(\mathbf{v}) = \int_{\Omega} \int_{\Omega} [V_0(x) - V_1(x + \mathbf{v}(z))]^2 B_{\mathcal{N}}(z - x) dx dz, \quad (14)$$

where $B_{\mathcal{N}}$ is a box filter of width $|\mathcal{N}|$. After discretization we arrive at the following cost for a single pixel p :

$$\text{SSD}(p, \mathbf{v}_p) = \sum_{q \in \mathcal{N}(p)} |V_0(q) - V_1(q + \mathbf{v}_p)|^2 \omega(q - p), \quad (15)$$

where we have set $\omega(q - p) = \frac{1}{|\mathcal{N}(p)|}$. A locally valid, convex approximation would be a 1st-order Taylor-expansion around $q + \mathbf{v}_p$ for all voxels $q \in \mathcal{N}(p)$. However, this requires, at each location, the computation of the gradient for all $|\mathcal{N}(p)|$ voxels in the neighborhood. A computationally more efficient idea is to expand (15) around the current flow estimate $\mathbf{v}_{0,q}$ for each voxel q . *I.e.*, after multiplying with λ we use as convexified data term:

$$\begin{aligned} G_{\text{SSD}}(p, \mathbf{v}_p) &= \lambda \sum_{q \in \mathcal{N}(p)} |V_0(q) - V_1(q + \mathbf{v}_{0,q})| \\ &\quad - (\mathbf{v}_p - \mathbf{v}_{0,q})^\top \nabla V_1|_{(q + \mathbf{v}_{0,q})}|^2 \omega(q - p). \end{aligned} \quad (16)$$

With this formulation the gradients and volumes must be evaluated only once per voxel, namely at the current flow estimate. The proximal map for the SSD at pixel p amounts to solving a small quadratic problem per voxel to update \mathbf{v}_p :

$$\left((I + \tau \partial G)^{-1} (\hat{\mathbf{v}}) \right)_p = \underset{\mathbf{v}_p}{\text{argmin}} \frac{1}{2\tau} (\hat{\mathbf{v}}_p - \mathbf{v}_p)^2 + G_{\text{SSD}}(p, \mathbf{v}_p). \quad (17)$$

Semi-Dense Flow. For large volumes the global optimization of (10) is both computationally expensive and memory-hungry. However, fluid flow is somewhat special: The need for very high image resolution arises mainly from the data term. Image gradients are present only along particle silhouettes, hence, the resolution must be chosen such that a voxel is at most $\frac{1}{2}$ particle diameter, to ensure individual particles are visible. Moreover, the PIV literature recommends particle diameters > 2 pixels in order to avoid peak locking effects (bias towards integer values) [30, 31]. On the contrary, by reconstructing the flow field only from the particles one makes an implicit assumption that each particle's motion is representative of the surrounding volume – patterns smaller than the spacing between two particles cannot be resolved for lack of evidence. The effective resolution of the flow is thus significantly lower.

We exploit that situation: Flow vectors are estimated at a lower grid resolution (per default we use a spacing of $h=4$ original voxels), whereas the data cost is nevertheless evaluated at full resolution, so as to preserve the particle boundaries implicit in the high-frequency gradients. If required, missing flow vectors can be interpolated from the sparser grid, without significant loss of accuracy. We do this in our evaluation. Note, besides computational savings the sparser flow grid has the additional advantage that the regularizer operates over larger, more meaningful spatial scales.

Local Flow Baseline. As a baseline, and to assess different data costs independent of the regularization scheme, we implement a simple local flow estimation similar to the local correlation schemes used in standard PIV. The baseline crops a 3D interrogation volume around each pixel of

the first volume, and exhaustively computes the similarity to the second volume at all possible voxel positions within ± 5 voxels per dimension. The best position is refined to sub-voxel accuracy by fitting a quadratic polynomial to the similarity scores in a $3 \times 3 \times 3$ neighborhood.

The local matching strategy, without any explicit regularization, requires very large interrogation volumes to guarantee a fairly unique particle distribution. *E.g.*, Elsinga *et al.* [11] suggest $41 \times 41 \times 41$ voxels, so that each volume contains on average 25 particles. Note, the overlap between these large interrogation volumes also reduces the effective resolution of the flow field considerably, therefore also local methods often compute only a semi-dense flow field [11].

4. Evaluation

For a quantitative evaluation of our fluid flow method we use data from the Johns Hopkins Turbulence Database (JHTDB) [20, 27], which provides a direct numerical simulation of isotropic turbulent flow in incompressible fluids. To separate the mutual influence of tomographic and 3D flow reconstruction, we show results for the full pipeline, including our simple particle reconstruction from images, and when computing the flow from noise-free, ground truth, particle volumes. Moreover, we analyze the influence of different cost functions, window sizes, particle densities and spacings of flow vectors and compare various regularizers. Additional evaluation results, also on experimental data, can be found in the supplemental material.

Setup. We follow the test setup of “test case D” of the 4th International PIV Challenge [19]. The challenge took place in conjunction with the 17th International Symposium on Applications of Laser Techniques to Fluid Mechanics, where participants from industry and academia handed in results of their approaches on provided input data. Unfortunately no ground truth data of the challenge is provided, such that we had to generate new data with similar specifications using existing flow-fields from the JHTDB. Following the guidelines in [19], we use the same discretization level as [19] to obtain a volume of $1024 \times 512 \times 352$ voxels and read out the ground truth flow fields at each voxel position. The average magnitude of the 3D displacements is 1.9 voxel units, with a maximum of 5.4 voxels. To generate input data, tracer particles were randomly sampled inside the volume and rendered to four symmetric camera views with viewing angles of $\pm 35^\circ$ w.r.t. the yz -plane of the volume, respectively $\pm 18^\circ$ w.r.t. the xz -plane. Images have 1500×800 pixels, with intrinsics chosen such that the area of one projected voxel matches approximately the pixel area, and have standard 8-bit intensity range. The amount of particles is chosen to yield a density of 0.1 particles per pixel in image space, corresponding to $\approx \frac{3}{10000}$ particles per 3D voxel, or 0.4 particles in a $11 \times 11 \times 11$ window. Par-

ticles are rendered with varying diameters up to a maximum of 3 pixels, and varying brightness. Figure 2 shows the rendered image for one of the four cameras.

Tomographic Reconstruction. We reconstruct the particle volume for each time-step with 5 MART iterations (+ anisotropic Gaussian smoothing after each iteration, see above) and run the 3D flow estimation on the raw MART outputs. *I.e.*, we do not attempt to reconstruct explicit 3D particle locations, but rather continue with a 3D dense map of occupancy scores. In contrast, the quality metrics for our tomographic reconstruction pipeline requires discrete particles. To obtain them, we threshold the occupancy scores at $I_{\min}=0.01$ and run non-maxima suppression with a $3 \times 3 \times 3$ kernel (*c.f.* [19]). To control the experiments for errors and ambiguities of the reconstruction step, we additionally generate noise-free occupancy volumes directly from the simulation, using trilinear interpolation to map non-integer particle coordinates to particle intensities.

As a quality metric for the static reconstruction part in isolation, Elsinga *et al.* [11] defined the quality factor

$$Q = \frac{\sum E(i) \cdot E^r(i)}{\sqrt{\sum E(i)^2 \cdot \sum E^r(i)^2}}, \quad (18)$$

where $E(i)$ and $E^r(i)$ are the voxel intensity values of the estimated reconstruction and the reference volume. A further common metric, the power ratio, is defined as

$$PR = \frac{N_T}{N_G} \left(\frac{\langle I_T \rangle}{\langle I_G \rangle} \right)^2, \quad (19)$$

with N_T, N_G the number of true and ghost particles, and $\langle I_T \rangle, \langle I_G \rangle$ their mean intensities. For details see [19].

With our basic re-implementation of MART, we reach a quality factor $Q=0.77$. Our method reconstructs 98% of the true particles, but generates approximately three times as many ghost particles ($3.32 \cdot N_T$), leading to a power ratio of $PR=13.1$. Note, most ghost particles have lower intensities than true particles (see Figure 4). The reconstruction algorithm was not the focus of our work and ranks in the middle of the field when compared to other participants of the PIV challenge [19]. The best, highly engineered methods achieve quality factors little below 1 and power ratios >100 . Nevertheless, even with this simple reconstruction front-end, our flow estimation is competitive with the best published methods – see below.

Volumetric Flow. In all tests we use coarse-to-fine estimation with a pyramid scale factor of 0.95, 8 pyramid levels, 20 warps and 30 inner iteration per pyramid level. Unless specified otherwise, we use SSD with an IV of size 11^3 as data-term, QRD_∞ for regularization, step-size $h=4$ in the semi-dense approach, and a particle density of 0.1 particles per pixel. We use the average endpoint error (AEE) and

average angular error (AAE) as error metrics for our evaluation. To illustrate the reconstruction quality, we visualize a single xy -slice ($z=60$) of the flow field’s u -component in Figure 5. Although a few small-scale details are lost, our method provides a fairly detailed picture of the flow.

Regularizers. In Table 1 we show results of our approach for different regularizers. In contrast to 2D flow, our volume setup is not a projection of the 3D world, there are no discontinuities due to occlusion boundaries. Modeling in the 3D domain allows us to utilize physical constraints, which would be much harder in a 2D projection. In our opinion, this is why simpler, but physically less accurate proxy constraints like piecewise constancy (TV) work better in 2D, but not for our volumetric fluid flow. Overall, quadratic gradient regularization combined with a hard constraint on the divergence of the flow field (QRD_∞) achieves the best results. Furthermore, we tested different values of α for the soft penalty variant QRD_α in Table 2, reporting also the average absolute divergence (AAD) of the resulting flow field.

	TV	QR	QRD_α	QRD_∞
AEE	0.462	0.408	0.371	0.369
AAE	12.664	11.541	10.553	10.478

Table 1. Average endpoint error (AEE) and average angular error (AAE) for different regularizers. Volume: $1024 \times 512 \times 352$.

α	0	1	12	32	64	128
AEE	0.408	0.400	0.380	0.373	0.371	0.375
AAD	0.023	0.023	0.012	0.010	0.009	0.010

Table 2. Average endpoint error (AEE) and average absolute divergence (AAD) for varying α .

Interrogation Volume Size. We further investigated different sizes of the interrogation volume IV (*i.e.* the 3D matching window used for data cost computation) and show results in Table 3. Our experiments suggest an IV size of 11^3 , which compares favorably with the large windows (41^3) needed by local methods. A bigger IV size of 13^3 or 15^3 leads to no significant improvement.

IV	7	9	11	13	15
AEE	0.406	0.380	0.369	0.364	0.365
AAE	11.523	10.783	10.478	10.351	10.408

Table 3. Reconstruction quality in dependence of the IV size.

Particle Density. In order to be comparable with [19], we chose a particle density of 0.1 particles per pixel (ppp) for our experiments. The performance for different densities is compared in Tab. 4. Best results are achieved with 0.075ppp, indicating that, at higher densities, ambiguities in the 3D reconstruction impair the flow estimation.

ppp	0.05	0.075	0.1	0.125	0.15
AEE	0.349	0.340	0.369	0.448	0.571
AAE	9.810	9.571	10.478	12.929	16.736

Table 4. Influence of the particle density per pixel (ppp) on the reconstruction quality.

Stepsize. To justify our semi-dense flow computation at reduced resolution, we compare the results with different spacing h for a smaller volume of $256 \times 256 \times 256$, for which computation at full resolution is still tractable. Table 5 confirms that discretizing the flow field at a bit lower resolution than the data term causes no loss in quality. In fact, $h=2$ gave the best results, even beating the full resolution $h=1$. This suggests that realistic flow fields are indeed locally smooth at the resolution of single particles, and longer-range regularization improves the result. For the full volume ($1024 \times 512 \times 352$) and a step-size of $h=2$ we get an AEE of 0.341 voxel. We point out that the underlying flow field was generated for the PIV competition as a realistic example to challenge state-of-the-art flow estimation. The complexity (respectively, smoothness) of the flow pattern is representative of real, relevant fluid dynamics experiments.

h	1	2	4	6
AEE	0.347	0.342	0.362	0.468
AAE	10.484	10.334	10.859	14.383

Table 5. Average endpoint error (AEE) and average angular error (AAE) for flow vector spacing. Volume: $256 \times 256 \times 256$.

Noise-free Input. Next, we test the flow algorithm alone, starting from a noise-free particle volume, see Table 6. Performance improves significantly, suggesting that the flow estimation would perform even better in conjunction with a more sophisticated 3D reconstruction, *e.g.* using data from multiple time-steps to discard ghost particles [25, 38].

ppp	0.05	0.075	0.1	0.125	0.15
AEE	0.298	0.253	0.227	0.210	0.199
AAE	8.301	7.064	6.326	5.860	5.538

Table 6. Flow estimation from noise-free particle volume for different particle densities.

Multiple Time-steps. Even though focusing on flow estimation between two time-steps we tried a simple extension to multiple time-steps. Following [13], the flow at the current time-step is initialized by the estimated fluid motion of the previous time-step. This allows one to skip coarser pyramid levels and compute the flow directly at the highest resolution. However, the quantitative gain in performance is marginal (AEE of 0.362 after 3 time-steps instead of 0.369 for two-frame result).

Comparison with other Methods. Since we do not have access to the ground truth from the PIV Challenge [19], or




Figure 5. xy -slice of the flow in X -direction. *left*: Ground truth. *center*: Estimated from noise-free particle distribution. *right*: Estimated from MART reconstruction.

to any of the participating methods, we unfortunately cannot directly compare to their results. However, we have generated our test setup according to the specifications of the challenge, using the same flow database. Hence, we believe that the results are roughly comparable. The best performing methods in the challenge are DLR, LaVision and TUD. All three use information from more than 2 time steps to improve both the filtering of the 3D reconstruction and the motion estimation. DLR uses a combined PIV and particle tracking approach termed “shake the box” [33]. In a post-processing step they penalize divergence and high spatial frequencies. LaVision, a heavily engineered commercial software, and TUD both use a form of motion tracking-enhanced MART for reconstruction [25]. For the actual 3D PIV processing both methods use a combination of iterative volume deformation per time-step [8] and fitting of second-order polynomial trajectories to corresponding particle patterns over multiple time-steps [21]. Apparently, the two competitors only differ in their choice of parameters. In the challenge, DLR achieved best results with ≈ 0.25 voxels endpoint error (at a reconstruction power ratio of >100), LaVision and TUD reach endpoint errors ≈ 0.32 voxels. All other challenge participants report $AEE > 0.45$ voxels.

Assuming that the datasets are indeed comparable, our result is the best for a 2-frame method, and close to LaVision and TUD, who both exploit temporal coherence and highly optimized particle reconstruction codes. Moreover, the results for the noise-free particle reconstruction suggest that with a more sophisticated 3D reconstruction front-end, even the DLR result is within reach with only 2 time steps. We reiterate that, while we did our best to match the specifications of the PIV challenge, there inevitably will be differences between the datasets and the comparison is not exact.

To establish a reasonable 2-frame baseline for our method on the exact same dataset, we compare it to the exhaustive local flow baseline described in Section 3.2. The results are displayed in Table 7 for different cost functions. We show results for both the whole pipeline with 3D particle reconstruction and a noise-free particle volume (*c.f.* Table 6). For efficiency we only estimate and evaluate the local flow on a sparse grid with 16 voxel spacing (this does not change the result of local exhaustive search).

As discussed, the local method requires large neighborhood sizes to compute the similarity. Thus, high-frequency

	SSD	NCC	SAD	Census	QRD $_{\infty}$
<i>reconst.</i>	0.457	0.438	0.483	2.455	0.369
<i>noise-free</i>	0.416	0.408	0.425	0.446	0.227

Table 7. Average endpoint error (AEE) for the local flow baseline. Best result of the proposed scheme (SSD+QRD $_{\infty}$) carried over from Table 1 and 6 for comparison.

variations of the flow cannot be recovered well. Best results are achieved with SSD and NCC as data term, using an IV size of 41^3 . Even so, all tested cost functions have end point errors >0.4 voxels. Especially for the noise-free particle volume, errors are significantly higher than for our divergence-free variational scheme.

5. Conclusion

We have presented a volumetric flow estimation method, primarily aimed at fluid flow reconstruction via particle imaging velocimetry. Technically, our method is a 3D version of the canonical model for variational optical flow, augmented with a physically based regularizer for incompressible fluids. To handle realistic, high-resolution PIV data, we exploit that the flow field inherently has lower resolution than the particle images and process the data at two different resolutions: High resolution for the data term, which depends on small tracer particles, and lower resolution for the flow vectors and the regularization. Our method delivers high-quality flow estimates that compete with the state-of-the-art, although it is based on a rather naive particle reconstruction step and uses only information from two consecutive time steps. These limitations directly determine our future work: The next steps are flow estimation over longer time intervals, in order to exploit additional temporal consistency constraints; and a tighter coupling between flow estimation and 3D reconstruction, so as to mitigate the errors in the current (precomputed and frozen) particle volumes.

On a general note, there is still room for improvement in particle imaging velocimetry. We believe that a closer cooperation between researchers from fluid dynamics and computer vision could significantly boost the development of future 3D-PIV systems.

Acknowledgements. This work was supported by ETH grant 29 14-1. Christoph Vogel acknowledges support from the ERC starting grant 640156, HOMOVIS. We thank Markus Holzner and Michal Havlena for discussions and help with the evaluation.

References

- [1] R. Adrian and J. Westerweel. *Particle Image Velocimetry*. Cambridge University Press, 2011.
- [2] L. Alvarez, C. Castaño, M. García, K. Krissian, L. Mazorra, A. Salgado, and J. Sánchez. A new energy-based method for 3d motion estimation of incompressible PIV flows. *Computer Vision and Image Understanding*, 113(7), 2009.
- [3] M. Bai, W. Luo, K. Kundu, and R. Urtasun. Exploiting semantic information and deep matching for optical flow. *ECCV 2016*.
- [4] C. Bailer, B. Taetz, and D. Stricker. Flow fields: Dense correspondence fields for highly accurate large displacement optical flow estimation. *ICCV 2015*.
- [5] A. Chambolle and T. Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1), 2011.
- [6] A. Cheminet, B. Leclaire, F. Champagnat, A. Plyer, R. Yegavian, and G. Le Besnerais. Accuracy assessment of a lucaskanade based correlation method for 3D PIV. In *Int'l Symp. Applications of Laser Techniques to Fluid Mechanics*, 2014.
- [7] Q. Chen and V. Koltun. Full flow: Optical flow estimation by global optimization over regular grids. *CVPR 2016*.
- [8] S. Discetti and T. Astarita. Fast 3D PIV with direct sparse cross-correlations. *Experiments in Fluids*, 53(5), 2012.
- [9] S. Discetti, A. Natale, and T. Astarita. Spatial filtering improved tomographic PIV. *Experiments in Fluids*, 54(4), 2013.
- [10] A. Dosovitskiy, P. Fischer, E. Ilg, P. Häusser, C. Hazırbaş, V. Golkov, P. v.d.Smagt, D. Cremers, and T. Brox. FlowNet: Learning optical flow with convolutional networks. *ICCV 2015*.
- [11] G. E. Elsinga, F. Scarano, B. Wieneke, and B. W. Oudheusden. Tomographic particle image velocimetry. *Experiments in Fluids*, 41(6), 2006.
- [12] S. Gesemann, F. Huhn, D. Schanz, and A. Schröder. From noisy particle tracks to velocity, acceleration and pressure fields using B-splines and penalties. *Int'l Symp. on Applications of Laser Techniques to Fluid Mechanics*, 2016.
- [13] J. Gregson, I. Ihrke, N. Thuerey, and W. Heidrich. From capture to simulation: connecting forward and inverse problems in fluids. *ACM Transactions on Graphics*, 33(4), 2014.
- [14] D. Heitz, E. Mémin, and C. Schnörr. Variational fluid flow measurements from image sequences: synopsis and perspectives. *Experiments in Fluids*, 48(3), 2010.
- [15] G. T. Herman and A. Lent. Iterative reconstruction algorithms. *Computers in Biology and Medicine*, 6(4), 1976.
- [16] S. Hermann and R. Werner. High accuracy optical flow for 3D medical image registration using the census cost function. *Pacific-Rim Symp. Image and Video Technology*, 2014.
- [17] B. K. P. Horn and B. G. Schunck. Determining optical flow. *Artificial Intelligence*, 17, 1981.
- [18] Y. Hu, R. Song, and Y. Li. Efficient coarse-to-fine patch match for large displacement optical flow. *CVPR 2016*.
- [19] C. J. Kähler, T. Astarita, P. P. Vlachos, J. Sakakibara, R. Hain, S. Discetti, R. La Foy, and C. Cierpka. Main results of the 4th International PIV Challenge. *Experiments in Fluids*, 57(6), 2016.
- [20] Y. Li, E. Perlman, M. Wan, Y. Yang, C. Meneveau, R. Burns, S. Chen, A. Szalay, and G. Eyink. A public turbulence database cluster and applications to study Lagrangian evolution of velocity increments in turbulence. *Journal of Turbulence*, 9, 2008.
- [21] K. Lynch and F. Scarano. A high-order time-accurate interrogation method for time-resolved PIV. *Measurement Science & Technology*, 24(3), 2013.
- [22] H. G. Maas, A. Gruen, and D. Papantoniou. Particle tracking velocimetry in three-dimensional flows – part 1: Photogrammetric determination of particle coordinates. *Experiments in Fluids*, 15(2), 1993.
- [23] T. Mansi, X. Pennec, M. Sermesant, H. Delingette, and N. Ayache. iLogDemons: A demons-based registration algorithm for tracking incompressible elastic biological tissues. *Int'l Journal of Computer Vision*, 92(1), 2011.
- [24] M. Menze, C. Heipke, and A. Geiger. Discrete optimization for optical flow. *GCPR 2015*.
- [25] M. Novara, K. J. Batenburg, and F. Scarano. Motion tracking-enhanced MART for tomographic PIV. *Measurement Science & Technology*, 21(3), 2010.
- [26] F. P. Oliveira and J. M. R. Tavares. Medical image registration: a review. *Computer Methods in Biomechanics and Biomedical Engineering*, 17(2), 2014.
- [27] E. Perlman, R. Burns, Y. Li, and C. Meneveau. Data exploration of turbulence simulations using a database cluster. In *ACM/IEEE Conf. on Supercomputing*, 2007.
- [28] T. Pock, D. Cremers, H. Bischof, and A. Chambolle. An algorithm for minimizing the Mumford-Shah functional. *ICCV 2009*.
- [29] T. Pock, M. Urschler, C. Zach, R. Beichel, and H. Bischof. A duality based algorithm for TV-L1 optical flow image registration. *MICCAI 2007*.
- [30] A. Prasad, R. Adrian, C. Landreth, and P. Offutt. Effect of resolution on the speed and accuracy of particle image velocimetry interrogation. *Experiments in Fluids*, 13(2), 1992.
- [31] M. Raffel, C. E. Willert, S. Wereley, and J. Kompenhans. *Particle image velocimetry: a practical guide*. Springer, 2013.
- [32] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *PhysD*, 60(1-4), 1992.
- [33] D. Schanz, S. Gesemann, and A. Schröder. Shake-The-Box: Lagrangian particle tracking at high particle image densities. *Experiments in Fluids*, 57(5), 2016.
- [34] J. F. Schneiders and F. Scarano. Dense velocity reconstruction from tomographic ptv with material derivatives. *Experiments in Fluids*, 57(9), 2016.
- [35] L. Sevilla-Lara, D. Sun, V. Jampani, and M. Black. Optical flow with semantic segmentation and localized layers. *CVPR 2016*.
- [36] F. Steinbrücker, T. Pock, and D. Cremers. Advanced data terms for variational optic flow estimation. *VMV 2009*.
- [37] C. Vogel, S. Roth, and K. Schindler. An evaluation of data costs for optical flow. *GCPR 2013*.
- [38] B. Wieneke. Iterative reconstruction of volumetric particle distribution. *Measurement Science & Technology*, 24(2), 2013.
- [39] C. Zach, T. Pock, and H. Bischof. A duality based approach for realtime TV-L1 optical flow. *DAGM 2007*.