Variational 3D-PIV for incompressible fluid flow estimation

Katrin Lasinger[1], Christoph Vogel[2] and Konrad Schindler[1]

[1] Institute of Photogrammetry and Remote Sensing, ETH Zurich, Zurich, Switzerland
katrin.lasinger@geod.baug.ethz.ch

[2] Institute of Computer Graphics and Vision, TU Graz, Graz, Austria

We present a method to estimate a 3D velocity field from 2D particle images captured at different time instants. For each individual time step our work follows the general setting of TomoPIV. We propose to rely on dense, variational flow estimation instead of local correlation matching to recover the velocity field from uncertain and noisy 3D particle volumes. We start from the standard variational 2D flow model widely used in computer vision [4], extend it to 3D and adapt it to the requirements of fluid flow estimation. To account for the physical properties of incompressible fluids, we strictly enforce a divergence-free flow with an energy derived from the stationary Stokes equation, and additionally penalise the squared gradient of the flow. To cope with the computational demands of large volumetric datasets, we propose a semi-dense formulation, where flow is estimated and regularised at a lower spatial resolution, while the data term is evaluated at full image/voxel resolution to preserve the geometric detail and discriminative power of the local particle distribution. We also perform a systematic evaluation of different similarity measures and regularisers. We find that the optic flow requires a lot smaller interrogation volumes. Furthermore, a simple sum of squared differences (SSD) works well as data term.

Methodology. Our approach consists of two major steps: (1) 3D reconstruction per time step and (2) variational flow estimation from two consecutive time steps. For the 3D reconstruction step we follow [3] and run 5 iterations of multiplicative algebraic reconstruction (MART). In our implementation, we add anisotropic Gaussian smoothing with a 3x3x1 voxel kernel after every MART iteration, to account for elongated particle reconstructions along the *z*-axis due to the camera setup [2]. We refrain from explicit reconstruction of individual particles and instead use the soft occupancy probabilities obtained by MART as input for the subsequent 3D flow estimation. We found it advantageous to apply non-linear contrast stretching to the raw MART scores before the subsequent flow computation, and transform the voxel intensities according to $V_{out} = (V_{in})^{0.7}$.
For the actual flow estimation from two consecutive MART volumes we set up a global, variational energy (respectively, negative log-likelihood) function of the form: $\lambda E_D(V_0, V_1, \vec{v}) + E_s(\vec{v})$, and minimise it with the primal-dual algorithm of Chambolle and Pock [1] to obtain the optimal flow field $\vec{v}$. The optimisation is embedded in a coarse-to-fine scheme to handle large displacements and to speed up the computation. We empirically tested different data cost functions E_D and regularisers E_S, among which the best results were achieved with SSD in a window of size 11x11x11 as data cost, and a linear combination of a quadratic regularisation of the flow field and a hard constraint on the flow field's divergence as regulariser:

$$\lambda \int_\Omega \int_\Omega \left|V_0(x) - V_1\big(x + \vec{v}(z)\big)\right|^2 B_{\mathcal{N}}(z-x)\,dx\,dz + \frac{1}{2}\int_\Omega |\nabla u|^2 + |\nabla v|^2 + |\nabla w|^2 + \delta_{\{0\}}(\nabla \cdot \vec{v})\,dx \to \min_{\vec{v}},$$

where $\vec{v} = (u, v, w)^T$ represents a mapping of points x in V_0 to points $(x + \vec{v})$ in V_1, $B_{\mathcal{N}}$ is a box filter of width $|\mathcal{N}|$ and δ_C the indicator function of the convex set *C*. Note that the regularisation encourages physical plausibility already when estimating the flow field, such that no further post-processing is needed.
For large volumes, the global optimisation is computationally expensive and memory-hungry. While we do not want to lose high-resolution details of the observed particle distribution, the effective spatial resolution of particle-based flow reconstruction is inherently limited, due to the relatively large spacing between particles: in a typical setup, there is on average less than 1 particle within a volume of 11x11x11 pixel units. Hence, we construct a semi-dense approach that operates on two different voxel grids. One with high (pixel-level) resolution for the data term, and one with lower resolution (e.g., $4\times$ lower in each dimension, respectively $64\times$ fewer variables) for the flow vectors and the regulariser.

Results. For quantitative evaluation, we use data from the Johns Hopkins Turbulence Database [6], which provides a DNS of isotropic turbulent flow in incompressible fluids. Particles are randomly located in a volume of 1024x512x352 voxels to yield a density of ≈0.1ppp, and rendered to four symmetric camera views of 1500x800 pixels each, with viewing angles of ±35° w.r.t. the *yz*-plane of the volume and ±18° w.r.t. the *xz*-plane. Particles are rendered with varying brightness and a maximum radius of 3 pixels. The average magnitude of the 3D displacements is 1.9 voxel units, with a maximum of 5.4 voxels.
To separate the influence of the MART reconstruction scheme from the one of the 3D flow computation, we evaluate them separately. Additionally, we show results both for the full pipeline including our (comparatively simple) particle

reconstruction from images, and for flow computation from noise-free particle volumes. Furthermore, we evaluate different spatial resolutions of the reconstructed flow vectors, and different cost functions.

With our basic re-implementation of MART, we are able to reconstruct 98% of the true particles, but also generate $3\times$ as many (mostly low-intensity) ghost particles, corresponding to a quality factor $Q=0.77$ (without further post-processing or multi-temporal smoothing). Results of the flow estimation are shown in Table 1 and Figure 1. In the table, we list average end point and angular errors of our variational flow approach. As baseline, we compare to standard normalised cross correlation (which requires much larger interrogation volumes of 41x41x41 voxels), followed by sub-voxel refinement by fitting a quadratic polynomial to the correlation scores. We additionally show results for a noise free particle volume where we directly use the original 3D positions of the simulated data as input. We find that the flow estimation delivers more correct velocity fields than simple local cross-correlation. Moreover, the results with noise-free particle volumes suggest that even better results are possible if one feeds the variational flow estimation with high-end particle reconstructions such as used in [5]. In Figure 2 we show qualitative results of our pipeline on test case D of the 4th International PIV Challenge [5], which used a particle volume of 4096x512x352 voxels (with 20vox/mm), and otherwise is similar to our own dataset.

Method	Average Endpoint Error	Average Angular Error
Local Cross-correlation from MART reconstruction	0.44 vox	12.83°
Local Cross-correlation from noise-free particle volume	0.41 vox	11.79°
Variational 3D Flow from MART reconstruction	0.37 vox	10.48°
Variational 3D Flow from noise-free particle volume	0.23 vox	6.33°

Table 1 Quantitative results for baseline implementation (local NCC) and our variational 3D flow approach.

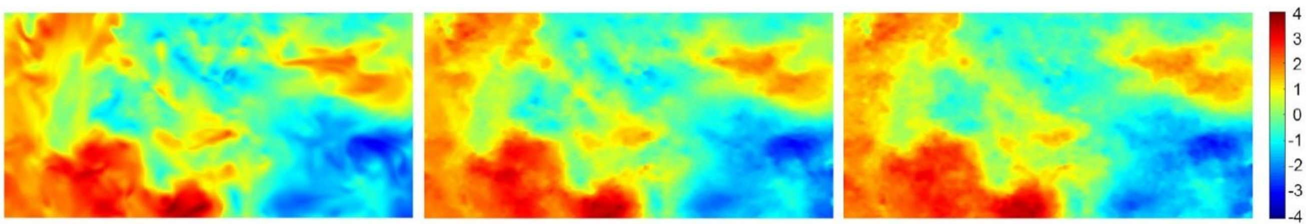

Figure 1 xy-slice of the flow in X-direction. *left*: ground truth. *center*: estimated from noise-free particle distribution. *right*: estimated from MART reconstruction.

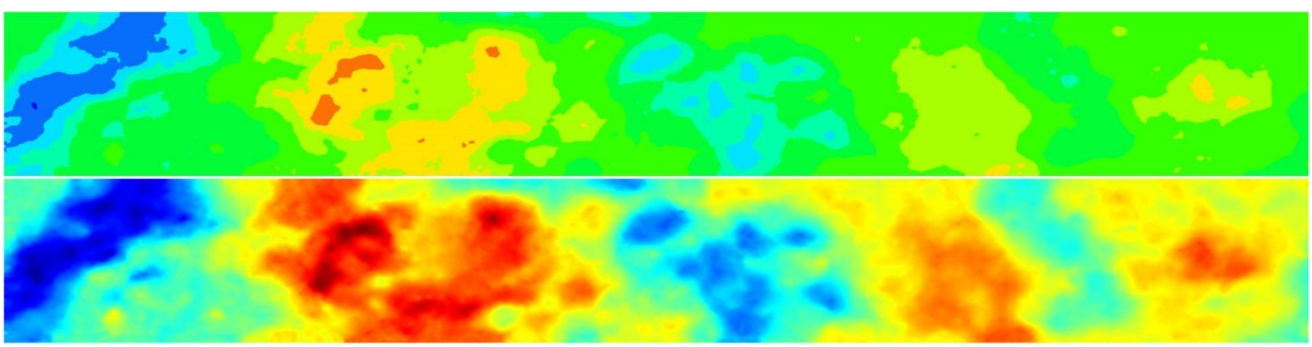

Figure 2 Reconstruction of snapshot 23 of case D of the 4th International PIV Challenge [5]. xy-slice of our estimated flow in X-direction on slice z=-5.2mm. *top*: same colour coding as in [5]. *bottom*: continuous colour coding.

REFERENCES

[1] Chambolle A, Pock T. “A first-order primal-dual algorithm for convex problems with applications to imaging” Journal of Mathematical Imaging and Vision 40-1 (2011) pp.120-45

[2] Discetti S, Natale A and Astarita T “Spatial filtering improved tomographic PIV” Experiments in Fluids 54.4 (2013) pp.1-13

[3] Elsinga GE, Scarano F, Wieneke B, and van Oudheusden BW “Tomographic particle image velocimetry” Experiments in Fluids 41.6 (2006) pp.933–947

[4] Horn BKP and Schunck BG “Determining optical flow” Artificial Intelligence 17 (1981) pp.185–203

[5] Kähler CJ, Astarita T, Vlachos PP, Sakakibara J, Hain R, Discetti S, Foy R and Cierpka C “Main results of the 4th International PIV Challenge” Experiments in Fluids 57.6 (2016) pp.1-71

[6] Li Y, Perlman E, Wan M, Yang Y, Meneveau C, Burns R, Chen S, Szalay A, Eyink G. “A public turbulence database cluster and applications to study Lagrangian evolution of velocity increments in turbulence” Journal of Turbulence 9 (2008) N31