A Joint Energy Formulation for 3D Particle Reconstruction and Velocity Field Estimation

Katrin Lasinger[1][*], Christoph Vogel[2], Thomas Pock[3,4], Konrad Schindler[1]

[1] ETH Zurich, Photogrammetry and Remote Sensing, Zurich, Switzerland

[2] Microsoft, Zurich, Switzerland

[3] TU Graz, Institute for Computer Graphics and Vision, Graz, Austria

[4] Austrian Institute of Technology, Vienna, Austria

[*] katrin.lasinger@geod.baug.ethz.ch

Abstract

We present a hybrid Lagrangian/Eulerian approach for 3D velocity field estimation from multi-view particle image sequences. By jointly optimizing over a sparse 3D particle representation and a dense velocity field we combine advantages of both 3D-PIV and 3D-PTV techniques. The sparse, explicit particle representation allows for memory-friendly, high-resolution particle reconstructions while the dense motion field representation facilitates incorporation of physical constraints. Our approach can handle high particle seeding densities in two-pulse setups. Furthermore, we present a simple extension to multiple time steps. Extra image information from an additional third time steps helps to resolve ambiguities in the particle reconstruction and allows for even higher particle seeding densities, up to 0.2ppp. We quantitatively evaluate our method on a synthetic dataset with varying seeding densities.

1 Introduction

Traditionally, there is a strict methodological distinction between Eulerian 3D-PIV and Lagrangian 3D-PTV for volumetric fluid velocity measurements from particle image sequences. Representative of the Eulerian formulation, Tomo-PIV (Elsinga et al., 2006) operates in voxel space and consists of two sequential steps: tomographic reconstruction and 3D velocity field estimation. While operating on a voxel grid allows for a straight-forward integration of physical priors (e.g., Heitz et al., 2010; **?**) and facilitates motion estimation also for two-pulse setups, major drawbacks of the method are the voxel-based particle representation and the necessity of large interrogation volumes for matching, leading to limited spatial resolution and large memory requirements. Furthermore, the sequential nature of the setup precludes the use of the subsequent motion estimates to improve particle locations and to remove of ghost particles. Lagrangian 3D-PTV methods such as Shake-the-Box (Schanz et al., 2016) directly reconstruct sparse 3D particle locations using iterative particle reconstruction (IPR, Wieneke, 2012) and track those particles over multiple time-steps. However, physical priors are only incorporated in a post-processing step, where the velocity field is interpolated onto a regular voxel grid.

In previous work (Lasinger et al., 2018) we have proposed a hybrid formulation that combines the memory-efficient, sparse representation of Lagrangian 3D particle locations with a dense, grid-based representation of the velocity field. The model can incorporate physical priors into the optimization of the motion field and, most importantly, makes it possible to jointly perform particle reconstruction and motion estimation, hence improving both. Specifically, we have derived a regularizer from the stationary Stokes equations, which is a reasonable approximation for a two-pulse setup (c.f., Ruhnau and Schnörr, 2007). In experiments with synthetic data it was quantitatively validated that this joint formulation supports higher particle seeding densities than traditional sequential approaches, by drastically reducing the amount of falsely reconstructed ghost particles. Thus, improving spatial resolution and accuracy of the estimated flow field.

Here we extend the joint approach to three consecutive time steps, in order to make further use of additional image information to better disambiguate particle locations and reduce ghost particles. I.e., for a
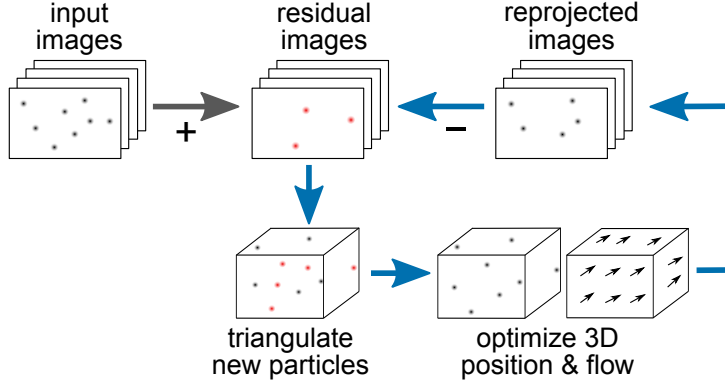
Figure 1: Our joint particle reconstruction and flow estimation pipeline: We alternate between triangulating new candidate particles from the residual images and jointly optimizing 3D particle positions and dense flow vectors. While triangulation is performed only for a single reference time step, the joint optimization can be performed over two or more consecutive time steps.

given time step we do not only consider the forward flow to the next time step but also the backward flow to a previous time step. This allows the handling of even higher particle seeding densities, as we will show in an extensive quantitative evaluation.

It is worth noting that, when considering multiple time steps, temporal coherence can be incorporated into the regularizer. Ruhnau et al. (2006) have derived an energy formulation from the incompressible vorticity transport equation for spatio-temporal regularization of 2D particle sequences. That energy is then minimized in an online fashion for consecutive image pairs. Besides penalizing the divergence of the estimated flow field, their energy term enforces smoothness between the estimated current vorticity and the vorticity field obtained from the previous image pair. Since vorticity is not known a-priori, the approach requires multiple initial time steps to converge.

For the scope of this paper we do not consider higher-order regularization but rather focus on the reduction of ghost particle generation, via joint optimization over multiple time steps. In Section 2 we recapitulate the joint optimization approach of Lasinger et al. (2018). Our extension to multiple time steps is presented in Section 3. Finally, we quantitatively evaluate our approach on synthetic data for varying seeding densities in Section 4 and show that information from additional time step helps to resolve ambiguities in the particle reconstruction.

2   Joint Optimization

We briefly summarize our 2-pulse approach for joint 3D particle reconstruction and flow estimation (Lasinger et al., 2018). An overview of the approach is given in Figure 1. The input to our method are images taken from three or more camera views observing the illuminated measurement volume of a fluid seeded with tracer-particles. We obtain the velocity field at a reference time step t_r from the estimated displacement field between the multi-view particle image sets obtained at t_r and t_{r+1}. To that end we alternate between generating new *candidate* 3D particles from triangulation at t_r, and a joint optimization of 3D particle positions and the dense motion field $\mathcal{U}_{r \to r+1}$, while enforcing sparsity of the particles set. Particles are represented at sparse 3D locations $\mathcal{P} := \{p_i\}_{i=1}^P$, $p_i \in \mathbb{R}^3$ with intensities $\mathcal{C} := \{c_i\}_{i=1}^P$, $c_i \in \mathbb{R}_0^+$, within the rectangular domain $\Omega \subset \mathbb{R}^3$. The motion field $\mathcal{U}_{r \to r+1}$ contains motion vectors $u \in \mathbb{R}^3$ at a finite set of positions $y \in \mathcal{Y} \subset \Omega$, which are arranged on a regular grid.

For particle proposal generation we follow a triangulation strategy similar to IPR (Wieneke, 2012). We start with a stricter triangulation threshold for the initial generation of putative particles and relax the threshold in later iterations. Triangulation is performed on the residuals between original camera images and synthetic images rendered from the currently estimated particle positions. The energy for the joint optimization consists of a data term E_D, a smoothness term E_S, operating on the motion field, and a sparsity term E_{Sp} operating on the intensities of the particles:

$$E(\mathcal{P}, \mathcal{C}, \mathcal{U}) := \frac{1}{2} E_D(\mathcal{P}, \mathcal{C}, \mathcal{U}) + \frac{\lambda}{2} E_S(\mathcal{U}) + \mu E_{Sp}(\mathcal{C}). \tag{1}$$

The data term E_D has the following form:

$$E_D(\mathcal{P}, \mathcal{C}, \mathcal{U}) := \sum_{k=1}^{K} \int_{\Gamma_k} \Big| I_k^{t_r}(x) - \sum_{i=1}^{P} \Pi_k(c_i \cdot \mathcal{N}(p_i, \sigma)(x)) \Big|_2^2 \mathrm{d}x +$$
$$\sum_{k=1}^{K} \int_{\Gamma_k} \Big| I_k^{t_{r+1}}(x) - \sum_{i=1}^{P} \Pi_k(c_i \cdot \mathcal{N}(p_i + u(p_i, \mathcal{U}_{r \to r+1}), \sigma)(x)) \Big|_2^2 \mathrm{d}x. \tag{2}$$

Following an additive (in terms of particles) image formation model, we integrate over the image plane Γ_k of camera k. $\Pi_k(.)$ denotes the projection from 3D to camera k, which can be implemented for both a pinhole and a polynomial camera model (Soloff et al., 1997). For the rendering we model particles (p_i, c_i) as Gaussian blobs with variance σ^2. Simply put, the data term penalizes differences between the original camera images at both time steps and corresponding images generated by re-projecting the estimated 3D particle locations. Note that particle triangulation is only performed for the first time step t_r. For rendering t_{r+1}, particles p_i are translated by the estimated motion field $\mathcal{U}_{r \to r+1}$.

The smoothness term E_S is derived from the stationary Stokes equations and enforces a divergence-free motion field as well as a quadratic regularization per component of the flow gradient (corresponding to the viscosity of the fluid):

$$E_S(\mathcal{U}) := \int_{\Omega} \sum_{l=1}^{3} |\nabla u_l(x, \mathcal{U})|_2^2 + \delta_{\{0\}}(\nabla \cdot u(x, \mathcal{U})) \mathrm{d}x. \tag{3}$$

The sparsity term E_{Sp} enforces sparsity of the reconstructed particle set, by suppressing (low-intensity) ghost particles. Furthermore, negative intensities are excluded from the set of feasible solutions:

$$E_{Sp}(\mathcal{C}) := \sum_{i=1}^{Q} |c_i|_0 + \delta_{\{\geq 0\}}(c_i). \tag{4}$$

To minimize energy (1) we employ the inertial variant of the proximal alternating linearized minimization method (PALM), which is suitable for our nonconvex and nonsmooth problem (Bolte et al., 2014; Pock and Sabach, 2016). The key idea of PALM is to split the variables into blocks, such that the problem is decomposed into one smooth function on the entire variable set, and a sum of non-smooth functions in which each block is treated separately. Furthermore, the optimization is embedded in a coarse-to-fine scheme, where we start with both a coarser grid $\mathcal{Y}$ for the dense motion field and a larger particle size for rendered particles to increase their area of influence. Both properties are refined in later iterations. We refer the reader to Lasinger et al. (2018) for further details about the method, as well as the choice of its free parameters.

3 Multiple Time Steps

In a multi-pulse setup, data from multiple time steps can be used to improve particle reconstruction and flow estimation. I.e., image data from additional time steps can help resolve ambiguities and reduce the number of ghost particles in the reconstruction. Moreover, temporal coherence can be enforced, e.g., by deriving an energy formulation from the full (incompressible) Navier Stokes equations. In this paper we focus only on improved 3D particle reconstruction and leave a possible extension to spatio-temporal regularization as future work. For a 2D version of the spatio-temporal regularizer, refer to Ruhnau (2006).

We extend the model of Lasinger et al. (2018) by an additional third time step. In addition to the estimated forward motion $\mathcal{U}_{r \to r+1}$ we estimate the backward motion from the reference time step to the previous time step t_{r-1}. Starting from t_r, we estimate the two motion fields $\mathcal{U}_{r \to r+1}$ and $\mathcal{U}_{r \to r-1}$ jointly, together with the 3D particle reconstruction at t_r. This requires adaptions of the data term and the smoothness term. Since particle triangulation is still performed only for the reference time step, the sparsity term remains unaltered. We schematically summarize the approach in Figure 2.

By extending (2) to three time steps we obtain the following data term:

$$E_D(\mathcal{P}, \mathcal{C}, \mathcal{U}) := \frac{1}{K} \sum_{j=r-1}^{r+1} \sum_{k=1}^{K} \int_{\Gamma_k} \Big| I_k^{t_j}(x) - \sum_{i=1}^{Q} \Pi_k(c_i \cdot \mathcal{N}(p_i + [j \neq r] \cdot u(p_i, \mathcal{U}_{r \to j}), \sigma)(x)) \Big|_2^2 \mathrm{d}x. \tag{5}$$
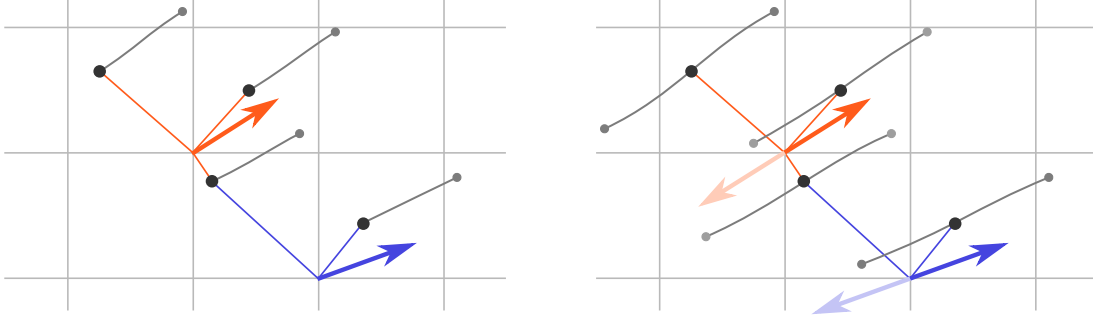
Figure 2: 2D Visualization of flow estimation at two selected grid locations (orange and blue vectors). The estimated flow field is obtained at discrete grid locations from data evidence at nearby particle locations (E_D) and a spatial regularization based on the stationary Stokes equations (E_S). *Left:* For two time steps, reconstructed particles at t_r are advected to the subsequent time step t_{r+1} with the estimated forward flow $\mathcal{U}_{r\to r+1}$. *Right:* In case of three time steps, particle locations are additionally advected by the estimated backward flow $\mathcal{U}_{r\to r-1}$ to coincide with input images from the previous time step. This significantly reduces the generation of ghost particles and leads to a more accurate flow estimation.

The smoothness term (3) is performed on both the forward flow $\mathcal{U}_{r\to r+1}$ and the backward flow $\mathcal{U}_{r\to r-1}$. Alternatively, one could average over the two flow fields and enforce (3) on the flow field $\mathcal{U}_r = \frac{1}{2}(\mathcal{U}_{r\to r+1} - \mathcal{U}_{r\to r-1})$.

4 Evaluation

To quantitatively compare our extension to the 2-pulse baseline and to evaluate the influence of different particle seeding densities, we generate synthetic particle tracks and ground truth velocity fields using the *Johns Hopkins Turbulence Database* (JHTDB, Li et al., 2008; Perlman et al., 2007). We follow the setup of test case D in the 4[th] International PIV Challenge (Kähler et al., 2016), where a symmetric four-camera setting was proposed with $\pm35°$ w.r.t. the *yz*-plane of the volume, respectively $\pm18°$ w.r.t. the *xz*-plane. However, we reduce the spatial extent in *x*-direction by a factor of four and use an image resolution of 1500×800 pixels. The ground truth velocity field is obtained at discrete DNS grid locations for a measurement volume of $256\times 128\times 88$ DNS points. Utilizing the proposed discretization level of the PIV challenge of 20 vox/mm the measurement volume correpsonds to $1024\times 512\times 352$ voxels. Note that the size of a voxel roughly corresponds to the size of a pixel in this setup. Particles are randomly sampled as floating point values within the z-range of 0-352. The x- and y-ranges were extended to both sides in order generate also particles outside the measurement region, reproducing experimental conditions of a thick laser illumnation. Particles are rendered to the camera images as Gaussian blobs with $\sigma = 1$ and varying intensity. We sampled data from the forced isotropic turbulence simulation of the JHTDB with a time difference of $\Delta t = 0.004$ between consecutive frames, resulting in a maximum flow magnitude of ≈ 7 voxels.

We compare our approach to two additional baseline methods. *IPR - sequential* independently performs sequential optimization of first the particle locations then the flow vectors. I.e., we separate (1) into particle reconstruction and subsequent motion field estimation. The second baseline (*HACKER*) isolates the flow estimation from the particle reconstruction. It starts from the *ground truth* 3D particle locations, hence no triangulation is required and no ghost particles occur. Since neither of the two baseline methods optimizes particle reconstruction and flow estimation together, their results for two and three time steps are identical.

In order to avoid triangulation and flow estimation issues at the border of the measurement volume, we extend the volume to both sides in x- and y-direction by 256 and 128 voxels respectively, resulting in $1536\times 768\times 352$ voxels. This allows to reconstruct also particles outside of the original measurement volume. For *HACKER* we chose to include only 3D particles that are visible in all four camera views. For quantitative evaluation only particles and flow estimates within the original measurement volume were considered.

Table 1 shows quantitative results of the four tested methods for seeding densities ranging from 0.075 to 0.225 particles per pixel (ppp). We measure the average endpoint error (EPE) between the estimated flow field $\mathcal{U}_{r\to r+1}$ and the ground truth velocity field at t_r in voxels. Furthermore, we state precision and recall

Table 1: Performance of different reconstruction methods under varying seeding density: our joint approach for two and three time steps, a sequential IPR approach, and velocity estimation from the ground truth particle locations. For the flow estimatoin error we measures the average endpoint error (EPE) in voxels, precision and recall of the particle reconstruction are stated in percentage (%).

| | IPR - joint | | | | | | IPR - sequential | | | HACKER |
| | 2 time steps | | | 3 time steps | | | | | | |
ppp	EPE	*prec.*	*recall*	EPE	*prec.*	*recall*	EPE	*prec.*	*recall*	EPE
0.075	0.174	90.23	99.86	0.163	93.19	99.87	0.169	91.46	99.20	0.148
0.1	0.149	89.14	99.79	0.140	91.50	99.84	0.171	70.99	96.20	0.132
0.125	0.138	84.14	99.70	0.126	89.70	99.79	0.235	32.55	74.41	0.122
0.15	0.125	77.00	99.29	0.116	83.48	99.67	0.304	22.59	53.46	0.114
0.175	0.122	72.34	98.00	0.110	75.33	99.30	0.361	17.58	40.35	0.110
0.2	0.122	61.54	96.00	0.108	75.42	98.28	0.393	14.17	34.28	0.108
0.225	0.128	47.53	92.59	0.111	61.94	96.26	0.444	11.57	27.44	0.106
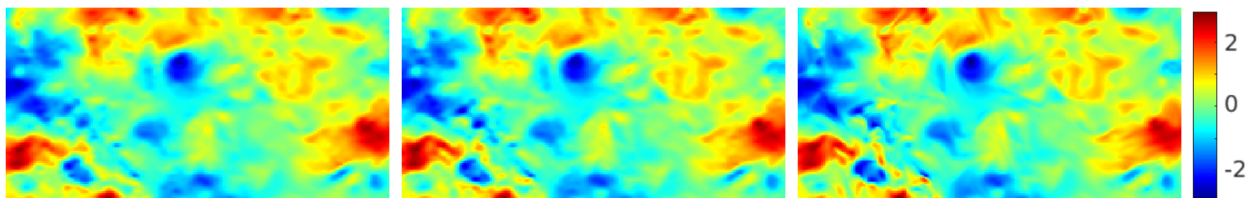
Figure 3: Visualization of xy-slice at z = 40 of the X-component of the flow estimated with our joint approach from 3 time steps. Left to right: 0.1ppp, 0.2ppp, ground truth. Note that with higher particle density details are recovered better.

of the reconstructed 3D particle set at t_r. The performance of HACKER serves as upper bound for the other methods. Based on a perfect particle set, it naturally improves with increasing seeding densities, due to the increased spatial evidence that allows for more accurate motion estimation. We observe that the sequential IPR approach starts to degrade already at relatively low seeding densities. Our joint approach handles higher seeding densities much better, resulting in both better particle reconstruction and flow estimation. The proposed extension to three time steps outperforms the two-time step variant in all tested seeding densities. Performance is on-par with *HACKER* up to 0.2ppp. Only at the highest tested seeding density of 0.225 the error of the estimated flow field starts to increase slightly. In Figure 3 we visualize the X-component of the flow field estimated with our joint 3-time-step approach for seeding densities of 0.1 and 0.2ppp togheter with the corresponding ground truth. Details at high frequency changes are recovered better with the higher seeding density.

5 Conclusion

We have presented a joint particle reconstruction and flow estimation approach for variational multi-camera 3D-PIV measurements. Our hybrid Eulerian/Lagrangian formulation makes it possible to handle high particle seeding densities, even in two-pulse setups. It also allows the use of the Stokes prior to obtain divergence-free motion fields. As a first step towards joint optimization over longer time sequences, we have extended the method to three time steps and demonstrated empirically that this already reduces ambiguities in the particle reconstruction and makes it possible to handle even higher seeding densities. Natural directions for future work are to use longer multi-pulse sequences, and to incorporate temporally coherent physical priors for multi-pulse sequences, e.g., by employing the full incompressible Navier-Stokes equations. While such upgrades will likely yield relatively small, diminishing returns in terms of the reconstructed velocity fields, they would allow one to extract additional properties of the fluid volume, such as for instance accelerations.

References

Bolte J, Sabach S, and Teboulle M (2014) Proximal alternating linearized minimization for nonconvex and nonsmooth problems. *Mathematical Programming* 146:459–494

Elsinga GE, Scarano F, Wieneke B, and van Oudheusden BW (2006) Tomographic particle image velocimetry. *Experiments in Fluids* 41:933–947

Heitz D, Mémin E, and Schnörr C (2010) Variational fluid flow measurements from image sequences: synopsis and perspectives. *Experiments in Fluids* 48:369–393

Kähler CJ, Astarita T, Vlachos PP, Sakakibara J, Hain R, Discetti S, La Foy R, and Cierpka C (2016) Main results of the 4th international piv challenge. *Experiments in Fluids* 57:97

Lasinger K, Vogel C, Pock T, and Schindler K (2018) 3d fluid flow estimation with integrated particle reconstruction. in *Pattern Recognition - 40th German Conference, GCPR 2018, Stuttgart, Germany, October 9-12, 2018, Proceedings*. pages 315–332

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

Perlman E, Burns R, Li Y, and Meneveau C (2007) Data exploration of turbulence simulations using a database cluster. in *SC '07: Proceedings of the 2007 ACM/IEEE Conference on Supercomputing*. pages 1–11

Pock T and Sabach S (2016) Inertial proximal alternating linearized minimization (iPALM) for nonconvex and nonsmooth problems. *SIAM Journal on Imaging Sciences* 9:1756–1787

Ruhnau P (2006) *Variational Fluid Motion Estimation with Physical Priors*. Ph.D. thesis. Universität Mannheim

Ruhnau P and Schnörr C (2007) Optical stokes flow estimation: an imaging-based control approach. *Experiments in Fluids* 42:61–78

Ruhnau P, Stahl A, and Schnörr C (2006) On-line variational estimation of dynamical fluid flows with physics-based spatio-temporal regularization. in *Pattern Recognition, 28th DAGM Symposium, Berlin, Germany, September 12-14, 2006, Proceedings*. pages 444–454

Schanz D, Gesemann S, and Schröder A (2016) Shake-the-box: Lagrangian particle tracking at high particle image densities. *Experiments in Fluids* 57:70

Soloff SM, Adrian RJ, and Liu ZC (1997) Distortion compensation for generalized stereoscopic particle image velocimetry. *Measurement Science and Technology* 8:1441–1454

Wieneke B (2012) Iterative reconstruction of volumetric particle distribution. *Measurement Science and Technology* 24:024008