

3D Fluid Flow Reconstruction Using Compact Light Field PIV

Zhong Li^{1*}, Yu Ji², Jingyi Yu^{2,3}, and Jinwei Ye⁴

¹ University of Delaware, Newark, DE, USA

² DGene, Baton Rouge, LA, USA

³ ShanghaiTech University, Shanghai, China

⁴ Louisiana State University, Baton Rouge, LA, USA
jye@csc.lsu.edu

Abstract. Particle Imaging Velocimetry (PIV) estimates the fluid flow by analyzing the motion of injected particles. The problem is challenging as the particles lie at different depths but have similar appearances. Tracking a large number of moving particles is particularly difficult due to the heavy occlusion. In this paper, we present a PIV solution that uses a compact lenslet-based light field camera to track dense particles floating in the fluid and reconstruct the 3D fluid flow. We exploit the focal symmetry property in the light field focal stacks for recovering the depths of similar-looking particles. We further develop a motion-constrained optical flow estimation algorithm by enforcing the local motion rigidity and the Navier-Stoke fluid constraint. Finally, the estimated particle motion trajectory is used to visualize the 3D fluid flow. Comprehensive experiments on both synthetic and real data show that using a compact light field camera, our technique can recover dense and accurate 3D fluid flow.

Keywords: volumetric flow reconstruction, particle imaging velocimetry (PIV), light field imaging, focal stack

1 Introduction

Recovering time-varying volumetric 3D fluid flow is a challenging problem. Successful solutions can benefit applications in many science and engineering fields, including oceanology, geophysics, biology, mechanical and environmental engineering. In experimental fluid dynamics, a standard methodology for measuring fluid flow is called Particle Imaging Velocimetry (PIV) [1]: the fluid is seeded with tracer particles, whose motions are assumed to follow the fluid dynamics faithfully, then the particles are tracked over time and their motion trajectories in 3D are used to represent the fluid flows.

Although being highly accurate, existing PIV solutions usually require complex and expensive equipment and the setups end up bulky. For example, standard laser-based PIV methods [16, 6] use ultra high speed laser beam to illuminate particles in order to track their motions. One limitation of these method is

* This work was performed when Zhong Li was a visiting student at LSU.

that the measured motion field only contains 2D in-plane movement restricted on the 2D fluid slice being scanned, as the laser beam can only scan one depth layer at a time. To fully characterize the fluid, it is necessary to recover the 3D flow motion within the entire fluid volume. Three-dimensional PIV such as tomographic PIV (Tomo-PIV) [9] use multiple cameras to capture the particles and resolve their depths in 3D using multi-view stereo. But such multi-camera systems need to be well calibrated and fully synchronized. More recently, the Rainbow PIV solutions [47, 46] use color to encode particles at different depths in order to recover the 3D fluid flow. However, this setup requires specialized illumination source with diffractive optics for color-encoding and the optical system needs to be precisely aligned.

In this paper, we present a flexible and low-cost 3D PIV solution that only uses one compact lenslet-based light field camera as the acquisition device. A light field camera, in essence, is a single-shot, multi-view imaging device [33]. The captured light field records 4D spatial and angular light rays scattered from the tracer particles. As commercial light field cameras (*e.g.* Lytro Illum and Raytrix R42) can capture high resolution light field, we are able to resolve dense particles in 3D fluid volume. Small baseline of the lenslet array further helps recover subtle particle motions at sub-pixel level. In particular, our method benefits from the post-capture refocusing capability of light field. We use the focal stack to establish correspondences among particles at different depths. To resolve heavily occluded particles, we exploit the focal stack symmetry (*i.e.*, intensities are symmetric in the focal stack around the ground truth disparity [25, 41]) for accurate particle 3D reconstruction.

Given the 3D locations of particles at each time frame, we develop a physics-based optical flow estimation algorithm to recover the particle’s 3D velocity field, which represents the 3D fluid flows. In particular, we introduce two new regularization terms to refine the classic variational optical flow [17]: 1) one-to-one particle correspondence term to maintain smooth and consistent flow motions across different time frames; and 2) divergence-free regularization term derived from the Navier-Stoke Equations to enforce the physical properties of incompressible fluid. These terms help resolve ambiguities in particle matching caused by similar appearances while enforcing the reconstruction to obey physical laws. Through synthetic and real experiments, we show that using a simple single camera setup, our approach outperforms state-of-the-art PIV solutions on recovering volumetric 3D fluid flows of various types.

2 Related Work

In computer vision and graphics, much effort has been made in modeling and recovering transparent objects or phenomena directly from images (*e.g.*, fluid [32, 49], gas flows [20, 4, 28, 48], smoke [12, 14], and flames [13, 19], *etc.*). As these objects do not have their own appearances, often a known pattern is assumed and the light paths traveled through the transparent medium are estimated for 3D reconstruction. A comprehensive survey can be found in [18]. However, many

of these imaging techniques are designed to recover the 3D density field, which does not explicitly reveal the internal flow motion.

Our method, instead, aims at estimating the 3D flow motion in terms of velocity field. The measurement procedure is similar to the Particle Imaging Velocimetry (PIV) method that estimates flow motion from movement by injecting tracer particles. Traditional PIV [16, 6] recovers 2D velocity fields on thin fluid slices using high speed laser scanning. As 3D volumetric flow is critical to fully characterize the fluid behavior, recovering a 3D velocity field within the entire volume is of great interest.

To recover 3D velocity field of a dense set of particles, stereoscopic cameras [3, 35] are used to estimate the particle depth. Tomographic PIV (Tomo-PIV) [9, 36, 22] use multiple (usually three to six) cameras to determine 3D particle locations by space carving. Aguirre-Pablo *et al.* [2] perform Tomo-PIV using mobile devices. However, the accuracy of reconstruction is compromised due to the low resolution of mobile cameras. Other notable 3D PIV approaches include defocusing PIV [45, 21], Holographic PIV [50, 39], and synthetic aperture PIV [5, 31]. All these systems use an array of cameras for acquisition and each measurement requires elaborate calibration and synchronization. In contrast, our setup is more flexible by using a single compact light field camera. Recently proposed rainbow PIV [46, 47] use color-coded illumination to recover depth from a single camera. However, both the light source and camera are customized with special optical elements and only sparse set of particles can be resolved. Proof-of-concept simulations [27] and experiments [10] using compact light field or plenoptic cameras for PIV have been performed and showed efficacy. However, the depth estimation and particle tracking algorithms used in these early works are rather primitive and are not optimized according to light field properties. As result, the recovered particles are relatively sparse and the reconstruction accuracy is lower than traditional PIV. Shi *et al.* [37, 38] use ray tracing to estimate particle velocity with a light field camera, and conduct comparison with Tomo-PIV. In our approach, we exploit the focal stack symmetry [25] of light fields for more accurate depth reconstruction in presence of heavily occluded dense particles.

To recover the flow motion, standard PIV uses 3D cross-correlation to match local windows between neighboring time frames [9, 44]. Although many improvements (for instance, matching with adaptive window sizes [22]) have been made, the window-based solutions suffer problems at regions with few visible particles. Another class of methods directly track the path of individual particles over time [29, 36]. However, with increased particle density, tracking is challenging under occlusions. Heitz *et al.* [15] propose the application of variational optical flow to fluid flow estimation. Vedula *et al.* [43] extend the optical flow to dynamic environment and introduce the scene flow. Lv *et al.* [26] use a neural network to recover 3D scene flow. Unlike natural scenes that have diverse features, our PIV scenes only contain similar-looking particles. Therefore, existing optical flow or scene flow algorithms are not directly applicable to our problem. Some methods [47, 23] incorporate physical constraints such as the Stokes equation into optical flow framework to recover fluid flows that obey physical




Fig. 1. Overall pipeline of our light field PIV 3D fluid flow reconstruction algorithm.

laws. However, these physics-based regularizations are in high-orders and are difficult to solve. In our approach, we introduce two novel regularization terms: 1) rigidity-enforced particle correspondence term and 2) divergence-free term to refine the basic variational optical flow framework for estimating the motion of dense particles.

3 Our Approach

Fig. 1 shows the algorithmic pipeline of volumetric 3D fluid flow reconstruction using light field PIV. For each time frame, we first detect particles in the light field sub-aperture images using the IDL particle detector [7]. We then estimate particle depths through a joint optimization that exploits light field properties. After we obtain 3D particle locations, we compare two consecutive frames to establish one-to-one particle correspondences and finally solve the 3D velocity field using a constrained optical flow.

3.1 3D Particle Reconstruction

We first describe our 3D particle reconstruction algorithm that exploits various properties of light field.

Focal Stack Symmetry. A focal stack is a sequence of images focused at different depth layers. Due to the post-capture refocusing capability, a focal stack can be synthesized from a light field by integrating captured light rays. Lin *et al.* [25] conduct symmetry analysis on focal stacks and show that *non-occluding* pixels in a focal stack exhibit symmetry along the focal dimension centered at the in-focus slice. In contrast, occluding boundary pixels exhibit local asymmetry as the outgoing rays are not originated from the same surface. Such property is called focal stack symmetry. As shown in Fig. 2, in a focal stack, a particle exhibits symmetric defocus effect centered at the in-focus slice. It’s also worth noting that occluded particles could be seen in the focal stack as the occluder becomes extremely out-of-focus. Utilizing the focal stack symmetry helps resolve heavily occluded particles and hence enhances the accuracy and robustness of particle depth estimation.

Given a particle light field, we synthesize a focal stack from the sub-aperture images by integrating rays from the same focal slice. Each focus slice f has a




Fig. 2. Focal stack symmetry. We show zoom-in views of four focal slices on the right. A particle exhibits symmetric defocus effect (e.g., 31.5mm and 36.5mm slices) centered at the in-focus slice (34mm). In the 39mm slice, an occluded particle could be seen as the occluder becomes extremely out-of-focus.

corresponding disparity d that indicates the in-focus depth layer. Let $I(p, f)$ be the intensity of a pixel p at focal slice f . For symmetry analysis, we define an in-focus score $\kappa(p, f)$ a pixel p at focal slice f as:

$$\kappa(p, f) = \int_0^{\delta_{max}} \rho(I(p, f + \delta) - I(p, f - \delta)) d\delta \quad (1)$$

where δ represents tiny disparity/focal shift and δ_{max} is maximum shift amount; $\rho(\nu) = 1 - e^{-|\nu|^2/(2\sigma^2)}$ is a robust distance function with σ controlling its sensitivity to noises. According to the focal stack symmetry, the intensity profile $I(p, f)$ is locally symmetric around the true surface depth. Therefore, if the pixel p is in focus at its true depth disparity \hat{d} , $\kappa(p, \hat{d})$ should be 0. Hence given an estimated disparity d at p , the closer distance between d and \hat{d} , the smaller the $\kappa(p, d)$. We then formulate the focal stack symmetry term β_{fs} for particle depth estimation by summing up $\kappa(p, d)$ for all pixels in a focal slice f with disparity d :

$$\beta_{fs}(d) = \sum_p \kappa(p, d) \quad (2)$$

Color and Gradient Consistency. Besides the focal stack symmetry, we also consider the color and gradient data consistency across sub-aperture images for depth estimation using data terms similar to [25]. Specifically, by comparing each sub-aperture image with the center view, we define a cost metric $C(i, p, d)$ as:

$$C(i, p, d) = |I_c(\omega(p)) - I_i(\omega(p + d(p)\chi(i)))| \quad (3)$$

where i is the sub-aperture image index; I_c and I_i refers to the center view and sub-aperture image respectively; $\omega(p)$ refers to a small local window centered around pixel p ; $d(p)$ is an estimate disparity at pixel p ; and $\chi(i)$ is a scalar that scale the disparity $d(p)$ according to the relative position between I_c and I_i as $d(p)$ is the pixel-shift between neighboring sub-aperture images.

The cost metric C measures the intensity similarity between shifted pixels in sub-aperture images given an estimated disparity. By summing up C for all pixels, we obtain the sum of absolute differences (SAD) term for color consistency measurement:

$$\beta_{sad}(d) = \frac{1}{N} \sum_{i \in N} \sum_p C \quad (4)$$

where N is the total number of sub-aperture images (excluding the center view).

Besides the color consistency, we also consider the consistency in gradient domain. We first take partial derivatives of cost metric C (Eq. 3) in both x and y directions: $D_x = \partial C / \partial x$ and $D_y = \partial C / \partial y$ and then formulate the following weighted sum of gradient differences (GRAD) for gradient consistency measurement:

$$\beta_{grad}(d) = \frac{1}{N} \sum_{i \in N} \sum_p \mathcal{W}(i) D_x + (1 - \mathcal{W}(i)) D_y \quad (5)$$

In Eq. 5, $\mathcal{W}(i)$ is a weighing factor that determines the contribution of horizontal gradient cost (D_x) according to the relative positions of the two sub-aperture images being compared. It is defined as $\mathcal{W}(i) = \frac{\Delta i_x}{\Delta i_x + \Delta i_y}$, where Δi_x and Δi_y are the position differences between sub-aperture images along x and y directions. For example, $\mathcal{W}(i) = 1$ if the target view is located at the horizontal extent of the reference view. In this case, only the gradient costs in the x direction are aggregated.

Particle Depth Estimation. Finally, combining Eq. 2, 4, and 5, we form the following energy function for optimizing the particle disparity d :

$$\beta(d) = \beta_{fs}(d) + \lambda_{sad} \beta_{sad}(d) + \lambda_{grad} \beta_{grad}(d) \quad (6)$$

In our experiments, the two weighting factors are set as $\lambda_{sad} = 0.8$ and $\lambda_{grad} = 0.9$. We use the Levenberg-Marquardt (LM) optimization to solve Eq. 6. Finally, using the calibrated light field camera intrinsic parameters, we are able to convert the particle disparity map to 3D particle location. The pipeline of our 3D particle reconstruction algorithm is shown in Fig. 3.




Fig. 3. Our 3D particle reconstruction algorithm pipeline.

3.2 Fluid Flow Reconstruction

After we reconstruct 3D particles in each frame, we compare two consecutive frames to estimate the volumetric 3D fluid flow.

Given two sets of particle locations S_1 and S_2 recovered from consecutive frames, we first convert S_1 and S_2 into voxelized 3D volumes as occupancy probabilities Θ_1 and Θ_2 through linear interpolation. Our goal is to solve per-voxel 3D velocity vector $\mathbf{u} = [u, v, w]$ for the whole volume.

In particular, we solve this problem under the variational optical flow framework [17] and propose two novel regularization terms, the correspondence term and the divergence-free term, for improved accuracy and efficiency. Our overall energy function E_{total} is combination of regularization terms and is written as:

$$E_{total} = E_{data} + \lambda_1 E_{smooth} + \lambda_2 E_{corres} + \lambda_3 E_{div} \quad (7)$$

where λ_1 , λ_2 , and λ_3 are term balancing factors. Please see our supplementary material for mathematical details of solving this energy function. In the following, we describe the algorithmic details of each regularization term.

Basic Optical Flow. The data term E_{data} and smooth term E_{smooth} are adopted from basic optical flow. They are derived from the brightness constancy assumption. E_{data} enforces consistency between occupancy possibilities Θ_1 and Θ_2 at corresponding voxels and E_{smooth} constrain the fluid motion to be piece-wise smooth. In our case, E_{data} and E_{smooth} can be written as:

$$E_{data}(\mathbf{u}) = \int \|\Theta_2(\mathbf{p} + \mathbf{u}) - \Theta_1(\mathbf{p})\|_2^2 d\mathbf{p} \quad (8)$$

$$E_{smooth}(\mathbf{u}) = \|\nabla \cdot \mathbf{u}\|_2^2 \quad (9)$$

where \mathbf{p} refers to a voxel in fluid volume and ∇ is the gradient operator.

Correspondence Term. We propose a novel correspondence term for more accurate flow estimation. Notice that E_{data} in the basic optical flow only enforces voxel-level consistency while particle-to-particle correspondences are not guaranteed. We therefore develop a correspondence term E_{corres} to enforce one-to-one particle matching. E_{corres} helps improve matching accuracy especially in regions with high particle density.

Let's consider two sets of particles: $S_1 = \{s_1 | s_1 \in \mathbb{R}^3\}$ as reference and $S_2 = \{s_2 | s_2 \in \mathbb{R}^3\}$ as target. E_{corres} enforces the one-to-one particle matching between the target and reference sets. To formulate E_{corres} , we first estimate correspondences between particles in S_1 and S_2 . We solve this problem by estimating transformations that map particles in S_1 to S_2 .

In particular, we employ a deformable graph similar to [42] that considers local geometric similarity and rigidity. To build the graph, we uniformly sample a set of particles in S_1 and use them as graph nodes $\mathbf{G} = \{g_1, g_2, g_3, \dots, g_m\}$. We

then aim to estimate a set of affine transformations $\mathbf{A} = \{A_i\}_{i=1}^m$ and $\mathbf{b} = \{b_i\}_{i=1}^m$ for each graph node. We then use these graph nodes as control points to deform particles in S_1 instead of computing transformations for individual particles. Given the graph node transformations \mathbf{A} and \mathbf{b} , we can transform every particle $s_1 \in S_1$ to its new location s'_1 using a weighted linear combination of graph nodes transformations:

$$s'_1 = f(s_1, \mathbf{A}, \mathbf{b}) = \sum_{i=1}^m \varpi_i(s_1) (\mathbf{A}(s_1 - g_i) + g_i + b_i) \quad (10)$$

where the weight $\varpi_i(s_1) = \max(0, (1 - \|s_1 - g_i\|^2/R^2)^3)$ models a graph node g_i influence on a particle $s_1 \in S_1$ according to their Euclidean distance. This restricts the particle transformation to be only affected by nearby graph nodes. In our experiment, we consider the nearest four graph nodes and R is the particle's distance to its nearest graph node.

To obtain the graph node transformations \mathbf{A} and \mathbf{b} , we solve an optimization problem with energy function:

$$\Psi_{total} = \Psi_{data} + \alpha_1 \Psi_{rigid} + \alpha_2 \Psi_{smooth} \quad (11)$$

Ψ_{data} is the data term aims to minimize particle-to-particle distances after transformation and is thus formulated as:

$$\Psi_{data} = \sum_{s_1 \in S_1} \|s'_1 - c_i\|^2 \quad (12)$$

where c_i is the closest point to s'_1 in S_2 .

Ψ_{rigid} is a rigidity regularization term that enforces the local rigidity of affine transformation. Ψ_{rigid} can be written as:

$$\Psi_{rigid} = \sum_{\mathbf{G}} \|A_i^T A_i - \mathbb{I}\|_F^2 + (\det(A_i) - 1)^2 \quad (13)$$

where \mathbb{I} is an identity matrix.

The last term Ψ_{smooth} enforces the spatial smoothness of nearby nodes and is written as:

$$\Psi_{smooth} = \sum_{\mathbf{G}} \sum_{k \in \Omega(i)} \|A_i(g_k - g_i) + g_i + b_i - (g_k + b_k)\|^2 \quad (14)$$

where $\Omega(i)$ refers to the set of nearest four neighbors of g_i .

The overall energy function Ψ_{total} can be optimized with an iterative Gauss-Newton algorithm and the affine transformations \mathbf{A} and \mathbf{b} are thus solved. In our experiment, we use $\alpha_1 = 50$ and $\alpha_2 = 10$ for Eq. 11.

By applying Eq. 11, we can transform every particle $s_1 \in S_1$ to its new location s'_1 using the graph nodes' transformations. We then find S_1 's corresponding set S_2^c in the target S_2 using a nearest neighbor search (ie, $s_2^c = \text{nnsearch}(s'_1, s_2)$).

After we establish the one-to-one correspondences between S_1 and S_2 , our correspondence term can be formulated based on the color consistency assumption as follow:

$$E_{corres}(\mathbf{u}, S_1, S_2^c) = \sum_{s_1 \in S_1, s_2^c \in S_2^c} \|s_2^c - (s_1 + \mathbf{u}(s_1))\|_2^2 \quad (15)$$

We show the effectiveness of the correspondence term by comparing the velocity field obtained with vs. without E_{corres} . The results are shown in Fig. 4. This comparison demonstrates that our correspondence term greatly improves matching accuracy and hence benefits flow reconstruction.




Fig. 4. Particle matching between source and target volumes with vs. without using the correspondence term E_{corres} . In our plots, green lines indicate correct correspondences and red lines indicate incorrect ones.

Divergence-Free Term. To enforce the physical properties of incompressible fluid, we add a divergence-free regularization term E_{div} to the optical flow framework. Based on the Navier-Stoke equations, fluid velocity \mathbf{u} can be split into into two distinct components: irrotational component ∇P and solenoidal component $\mathbf{u}_{sol} = [u_{sol}, v_{sol}, w_{sol}]$ with the Helmholtz decomposition. The Irrotational component ∇P is curl-free and is determined by the gradient of a scalar function P (eg, pressure). The solenoidal component \mathbf{u}_{sol} is divergence-free and models an incompressible flow. From the divergence-free property, we have:

$$\nabla \cdot \mathbf{u}_{sol} = 0 \quad (16)$$

where $\nabla = [\frac{\partial}{\partial x}, \frac{\partial}{\partial y}, \frac{\partial}{\partial z}]^T$ is the divergence operator. Since $\mathbf{u} = \mathbf{u}_{sol} + \nabla P$, taking divergence on both sides, we have:

$$\nabla \cdot \mathbf{u} = \nabla^2 P \quad (17)$$

We solve Eq. 17 by Poisson integration and compute the scalar field as $P = (\nabla^2)^{-1}(\nabla \cdot \mathbf{u})$. We then project \mathbf{u} into the divergence-free vector field: $\mathbf{u}_{sol} = \mathbf{u} - \nabla P$. Similar to [11], we formulate a divergence-free term E_{div} that enforces the flow velocity field \mathbf{u} close to its divergence-free component \mathbf{u}_{sol} :

$$E_{div}(\mathbf{u}) = \|\mathbf{u} - \mathbf{u}_{sol}\|_2^2 \quad (18)$$

4 Experimental Results

To evaluate our fluid flow reconstruction algorithm, we perform experiments on both synthetic and real data under the light field PIV setting. We also evaluate our method on the John Hopkins Turbulence Database (JHUTDB) [24, 34] that has the ground truth fluid flow. All experiments are performed on a PC with Intel i7-4700K CPU with 16G of memory. On the computational time, the entire process takes about 2 minutes: 30 seconds for particle location estimation and 40 seconds for correspondence matching, and 50 seconds for velocity field reconstruction.

4.1 Synthetic Data

We first evaluate our proposed approach on simulated flows: a vortex flow and a drop flow. The flows are simulated within a volume of $100 \times 100 \times 20$ voxels. We randomly sample tracer particles within the fluid volume. The particle density is 0.02 per voxel. We render light fields images with angular resolution 7×7 and spatial resolution 434×625 . We simulate the advection of particles over time following the method in [40]. We apply our algorithms on the rendered light fields to recover 3D fluid flows. In Fig. 5, we show our recovered velocity fields in comparison with the ground truth ones. Qualitatively, our reconstructed vector fields are highly consistent with the ground truth ones.




Fig. 5. Synthetic results in comparison with the ground truth.

We perform quantitative evaluations using two error metrics: the average end-point error (AEE) and the average angular error (AAE). AEE is computed as the averaged Euclidean distance between the estimated particle positions and ground truth ones. AAE is computed with the average difference of vector in the velocity field. We compare our method with the multi-scale Horn-Schunck (H & S) [30] and the rainbow PIV [47]. Specifically, we apply H & S on our recovered 3D particles and use it as the baseline algorithm for flow estimation. With this comparison, we hope to demonstrate the effectiveness of our regularization terms in flow estimation. For rainbow PIV, we have implemented a renderer to generate depth-dependent spectral images of virtual particles. To ensure fairness, the

rendered images have the same spatial resolution as our input light field (ie, 434×625).

We also perform ablation study by testing two variants of our method: “w/o E_{corres} ” that takes out the correspondence term and “w/o E_{div} ” that takes out the divergence-free term. The experiments are performed on the vortex flow with particle density 0.02. Quantitative evaluations are shown in Fig. 6. The error maps of recovered velocity fields for our ablation study are shown Fig. 7. We can see that our method achieves the best performance when both regularization terms are imposed. Our outperforms both H & S and the rainbow PIV at various particle density levels. Further, our accumulated error over time grows much slower than the other two state-of-the-arts.




Fig. 6. Quantitative evaluation. The left two plots show errors with respect to different particle densities. The right two plots show accumulated errors over time.




Fig. 7. Ablation study. We show the error maps of estimated velocity field at three fluid volume slices.

4.2 John Hopkins Turbulence Database (JHUTDB)

Next we conduct experiments on data generated from the Johns Hopkins Turbulence Database (JHUTDB) [24]. To reduce processing time, we crop out a volume of $256 \times 128 \times 80$ voxels for each turbulence in the dataset. The norm of the velocity field at each location ranges from 0 to 2.7 voxels per time step. We generate random tracer particles with density 0.025 per voxels and advect the particles according to the turbulence velocity field. In our evaluation, we render

two light field images at two consecutive frames to estimate the particle locations and reconstruct the velocity field. Our reconstruction results in comparison with the ground truth is shown in Fig. 8. We show our reconstructed velocity volume in x, y, z directions. We also show the error map of magnitudes to illustrate that our method is highly accurate.

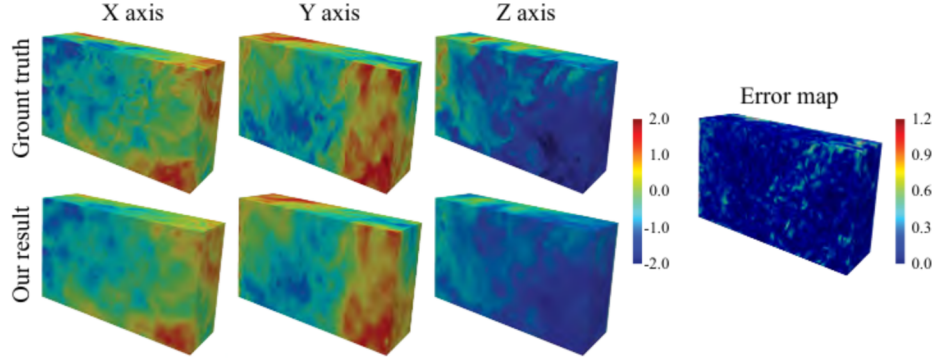


Fig. 8. JHUTDB velocity field reconstruction results.

4.3 Real Data

We finally test our method on real captured flow data. Fig. 9 shows our acquisition system for capturing real 3D flows. We use a Lytro Illum light field camera with 30mm focal length to capture the tracer particles in fluid. As Illum does not have video mode, we use an external control board to trigger the camera at high frequency to capture consecutive time frames. Due to the limitation of on-chip image buffer size, our acquisition cannot achieve very high frame rate. In our experiment, we set the trigger frequency to be 10 Hz. The capture light field has angular resolution 15×15 and spatial resolution 625×434 . We use the light field calibration toolbox [8] to process and decode raw light field data into sub-aperture images. We use the center view as reference for depth estimation and the effective depth volume that we are able to reconstruct is around




Fig. 9. Our real experiment setup. We use a compact light field camera in PIV setting.

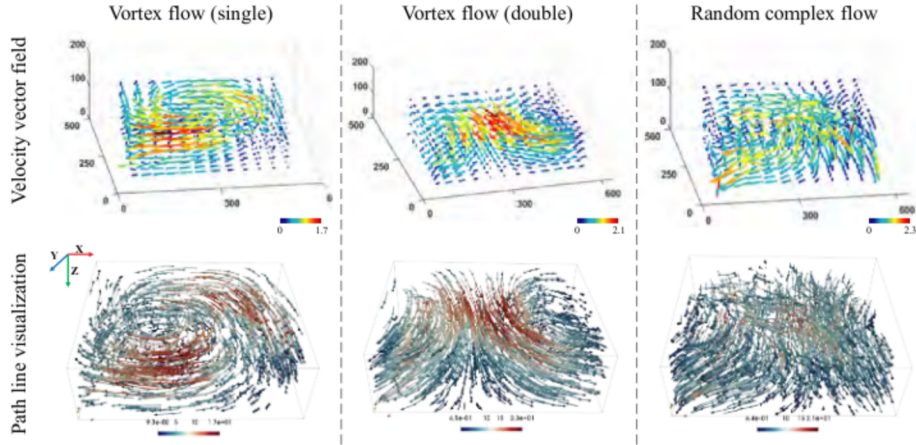


Fig. 10. Real experiment results. We show our recovered velocity fields (upper row) and path line visualizations on four consecutive frames (lower row) for three types of flows: vortex, double vortex and a random complex flow.

$600 \times 500 \times 200$ (mm), slightly lower than the capture image because we enforce rectangular volumes inside the perspective view frustum.

We use green polyethylene microspheres with density 1g/cc and size $1000\text{-}1180 \mu\text{m}$ as tracer particles. Before dispersing the particles, we mix some surfactant with the particles to reduce surface tension caused by water in order to minimize agglomeration between particles. We test on three types of flows: vortex, double vortex, and random complex flows.

Fig. 10 shows our recovered fluid flow velocity field and path line visualization (please refer to the supplemental material for more reconstruction results). We show three flow types, vortex, double vortex, and random complex flows. The left column shows the velocity field between first and the second frame. The right column shows the path line visualization through 1 - 4 frames. We can see that our reconstructions well depicts the intended fluid motions and are highly reliable.

We also compare our method with a recent state-of-the-art scene flow method [26] on the real data. The scene flow method takes two consecutive RGB-D images as inputs and use rigidity transform network and flow network for motion estimation. Since the method also needs depth map as input, we first calculate a depth map for the center view of light field and then combine the depth map with the sub-aperture color image and use them as input for [26]. The flow estimation results are shown in Fig. 11. We show the projected scene flows and the flow vector fields for three types of flows (single vortex, double vortex, and random flow). The scene flow method fails to recover the flow structures, especially for vortex flows. This is because our particles are heavily occluded and have very similar appearances. Further, the scene flow algorithm does not take the physical properties of fluid into consideration.

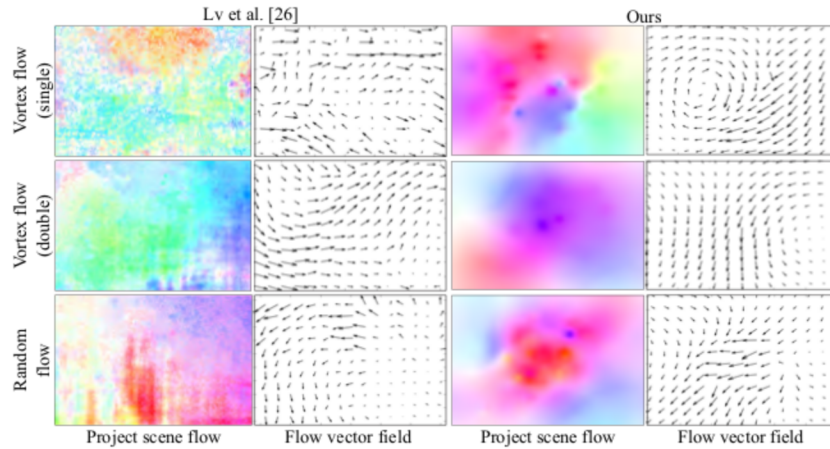


Fig. 11. Comparison result with scene flow (Lv et al. [26]) on the real data. We compare the project scene flow and the flow vector field on three types of flows.

5 Conclusions

In this paper, we have presented a light field PIV solution that uses a commercial compact light field camera to recover volumetric 3D fluid motion from tracer particles. We have developed a 3D particle reconstruction algorithm by exploiting the light field focal stack symmetry in order to handle heavily occluded particles. To recover the fluid flow, we have refined the classical optical flow framework by introducing two novel regularization terms: 1) the correspondence term to enforce one-to-one particle matching; and 2) the divergence-free term to enforce the physical properties of incompressible fluid. Comprehensive synthetic and real experiments as well as comparisons with the state-of-the-arts have demonstrated the effectiveness of our method.

Although our method can faithfully recover fluid flows in a small to medium volume, our method still has several limitations. First of all, due to the small baseline of compact light field camera, the resolvable depth range is rather limited. As a result, our volumetric velocity field’s resolution along the z-axis is much smaller than its x- or y-resolutions. One way to enhance the z-resolution is using a second light field camera capturing the fluid volume from an orthogonal angle. Second, in our fluid flow reconstruction step, only two consecutive frames are considered. Hence motion continuity might not always be satisfied. Adding temporal constraint to our optimization framework can be further improved.

Acknowledgements

This work is partially supported by the National Science Foundation (NSF) under Grant CBET-1706130 and CRII-1948524, and the Louisiana Board of Regent under Grant LEQSF(2018-21)-RD-A-10.

References

1. Adrian, R.J., Westerweel, J.: Particle image velocimetry. No. 30, Cambridge University Press (2011)
2. Aguirre-Pablo, A.A., Alarfaj, M.K., Li, E.Q., Hernández-Sánchez, J.F., Thoroddsen, S.T.: Tomographic particle image velocimetry using smartphones and colored shadows. In: Scientific Reports (2017)
3. Arroyo, M., Greated, C.: Stereoscopic particle image velocimetry. *Measurement science and technology* **2**(12), 1181 (1991)
4. Atcheson, B., Ihrke, I., Heidrich, W., Tevs, A., Bradley, D., Magnor, M., Seidel, H.P.: Time-resolved 3d capture of non-stationary gas flows. In: *ACM transactions on graphics (TOG)*. vol. 27, p. 132. ACM (2008)
5. Belden, J., Truscott, T.T., Axiak, M.C., Techet, A.H.: Three-dimensional synthetic aperture particle image velocimetry. *Measurement Science and Technology* **21**(12), 125403 (2010)
6. Brücker, C.: 3d scanning piv applied to an air flow in a motored engine using digital high-speed video. *Measurement Science and Technology* **8**(12), 1480 (1997)
7. Crocker, J.C., Grier, D.G.: Methods of digital video microscopy for colloidal studies. *Journal of colloid and interface science* **179**(1), 298–310 (1996)
8. Dansereau, D.G., Pizarro, O., Williams, S.B.: Decoding, calibration and rectification for lenselet-based plenoptic cameras. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1027–1034 (2013)
9. Elsinga, G.E., Scarano, F., Wieneke, B., van Oudheusden, B.W.: Tomographic particle image velocimetry. *Experiments in fluids* **41**(6), 933–947 (2006)
10. Fahringer, T., Thurow, B.: Tomographic reconstruction of a 3-d flow field using a plenoptic camera. In: *42nd AIAA Fluid Dynamics Conference and Exhibit*. p. 2826 (2012)
11. Gregson, J., Ihrke, I., Thuerey, N., Heidrich, W.: From capture to simulation: connecting forward and inverse problems in fluids. *ACM Transactions on Graphics (TOG)* **33**(4), 139 (2014)
12. Gu, J., Nayar, S.K., Grinspun, E., Belhumeur, P.N., Ramamoorthi, R.: Compressive structured light for recovering inhomogeneous participating media. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **35**, 1–1 (2013)
13. Hasinoff, S.W., Kutulakos, K.N.: Photo-consistent reconstruction of semitransparent scenes by density-sheet decomposition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **29**, 870–885 (2007)
14. Hawkins, T., Einarsson, P., Debevec, P.: Acquisition of time-varying participating media. In: *ACM Transactions on Graphics (ToG)*. vol. 24, pp. 812–815. ACM (2005)
15. Heitz, D., Mémin, E., Schnörr, C.: Variational fluid flow measurements from image sequences: synopsis and perspectives. *Experiments in fluids* **48**(3), 369–393 (2010)
16. Hori, T., Sakakibara, J.: High-speed scanning stereoscopic piv for 3d vorticity measurement in liquids. *Measurement Science and Technology* **15**(6), 1067 (2004)
17. Horn, B.K., Schunck, B.G.: Determining optical flow. *Artificial intelligence* **17**(1-3), 185–203 (1981)
18. Ihrke, I., Kutulakos, K.N., Lensch, H.P., Magnor, M., Heidrich, W.: Transparent and specular object reconstruction. In: *Computer Graphics Forum*. vol. 29, pp. 2400–2426. Wiley Online Library (2010)
19. Ihrke, I., Magnor, M.A.: Image-based tomographic reconstruction of flames. In: *Symposium on Computer Animation* (2004)

20. Ji, Y., Ye, J., Yu, J.: Reconstructing gas flows using light-path approximation. 2013 IEEE Conference on Computer Vision and Pattern Recognition pp. 2507–2514 (2013)
21. Kajitani, L., Dabiri, D.: A full three-dimensional characterization of defocusing digital particle image velocimetry. *Measurement Science and Technology* **16**(3), 790 (2005)
22. Lasinger, K., Vogel, C., Schindler, K.: Volumetric flow estimation for incompressible fluids using the stationary stokes equations. In: 2017 IEEE International Conference on Computer Vision (ICCV). pp. 2584–2592. IEEE (2017)
23. Lasinger, K., Vogel, C., Schindler, K.: Volumetric flow estimation for incompressible fluids using the stationary stokes equations. 2017 IEEE International Conference on Computer Vision (ICCV) pp. 2584–2592 (2017)
24. 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), N31 (2008)
25. Lin, H., Chen, C., Bing Kang, S., Yu, J.: Depth recovery from light field using focal stack symmetry. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 3451–3459 (2015)
26. Lv, Z., Kim, K., Troccoli, A., Sun, D., Rehg, J.M., Kautz, J.: Learning rigidity in dynamic scenes with a moving camera for 3d motion field estimation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 468–484 (2018)
27. Lynch, K., Fahringer, T., Thurow, B.: Three-dimensional particle image velocimetry using a plenoptic camera. In: 50th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition. p. 1056 (2012)
28. Ma, C., Lin, X., Suo, J., Dai, Q., Wetzstein, G.: Transparent object reconstruction via coded transport of intensity. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3238–3245 (2014)
29. Maas, H., Gruen, A., Papantoniou, D.: Particle tracking velocimetry in three-dimensional flows. *Experiments in Fluids* **15**(2), 133–146 (1993)
30. Meinhardt, E., Pérez, J.S., Kondermann, D.: Horn-schunck optical flow with a multi-scale strategy. *IPOL Journal* **3**, 151–172 (2013)
31. Mendelson, L., Techet, A.H.: Quantitative wake analysis of a freely swimming fish using 3d synthetic aperture piv. *Experiments in Fluids* **56**(7), 135 (2015)
32. Morris, N.J., Kutulakos, K.N.: Dynamic refraction stereo. *IEEE transactions on pattern analysis and machine intelligence* **33**(8), 1518–1531 (2011)
33. Ng, R., Levoy, M., Brédif, M., Duval, G., Horowitz, M., Hanrahan, P., et al.: Light field photography with a hand-held plenoptic camera. *Computer Science Technical Report CSTR* **2**(11), 1–11 (2005)
34. Perlman, E., Burns, R., Li, Y., Meneveau, C.: Data exploration of turbulence simulations using a database cluster. In: Proceedings of the 2007 ACM/IEEE conference on Supercomputing. p. 23. ACM (2007)
35. Pick, S., Lehmann, F.O.: Stereoscopic piv on multiple color-coded light sheets and its application to axial flow in flapping robotic insect wings. *Experiments in fluids* **47**(6), 1009 (2009)
36. Schanz, D., Gesemann, S., Schröder, A.: Shake-the-box: Lagrangian particle tracking at high particle image densities. *Experiments in fluids* **57**(5), 70 (2016)
37. Shi, S., Ding, J., Atkinson, C., Soria, J., New, T.H.: A detailed comparison of single-camera light-field piv and tomographic piv. *Experiments in Fluids* **59**, 1–13 (2018)

38. Shi, S., Ding, J., New, T.H., Soria, J.: Light-field camera-based 3d volumetric particle image velocimetry with dense ray tracing reconstruction technique. *Experiments in Fluids* **58**, 1–16 (2017)
39. Soria, J., Atkinson, C.: Towards 3c-3d digital holographic fluid velocity vector field measurement? tomographic digital holographic piv (tomo-hpiv). *Measurement science and technology* **19**(7), 074002 (2008)
40. Stam, J.: Stable fluids. In: *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*. pp. 121–128. ACM Press/Addison-Wesley Publishing Co. (1999)
41. Strecke, M., Alperovich, A., Goldluecke, B.: Accurate depth and normal maps from occlusion-aware focal stack symmetry. In: *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*. pp. 2529–2537. IEEE (2017)
42. Sumner, R.W., Schmid, J., Pauly, M.: Embedded deformation for shape manipulation. In: *ACM Transactions on Graphics (TOG)*. vol. 26, p. 80. ACM (2007)
43. Vedula, S., Baker, S., Rander, P., Collins, R.T., Kanade, T.: Three-dimensional scene flow. *Proceedings of the Seventh IEEE International Conference on Computer Vision* **2**, 722–729 vol.2 (1999)
44. Wieneke, B.: Volume self-calibration for 3d particle image velocimetry. *Experiments in fluids* **45**(4), 549–556 (2008)
45. Willert, C., Gharib, M.: Three-dimensional particle imaging with a single camera. *Experiments in Fluids* **12**(6), 353–358 (1992)
46. Xiong, J., Fu, Q., Idoughi, R., Heidrich, W.: Reconfigurable rainbow piv for 3d flow measurement. In: *Computational Photography (ICCP), 2018 IEEE International Conference on*. pp. 1–9. IEEE (2018)
47. Xiong, J., Idoughi, R., Aguirre-Pablo, A.A., Aljedaani, A.B., Dun, X., Fu, Q., Thoroddsen, S.T., Heidrich, W.: Rainbow particle imaging velocimetry for dense 3d fluid velocity imaging. *ACM Transactions on Graphics (TOG)* **36**(4), 36 (2017)
48. Xue, T., Rubinstein, M., Wadhwa, N., Levin, A., Durand, F., Freeman, W.T.: Refraction wiggles for measuring fluid depth and velocity from video. In: *European Conference on Computer Vision*. pp. 767–782. Springer (2014)
49. Ye, J., Ji, Y., Li, F., Yu, J.: Angular domain reconstruction of dynamic 3d fluid surfaces. In: *2012 IEEE Conference on Computer Vision and Pattern Recognition*. pp. 310–317. IEEE (2012)
50. Zhang, J., Tao, B., Katz, J.: Turbulent flow measurement in a square duct with hybrid holographic piv. *Experiments in Fluids* **23**(5), 373–381 (1997)