

Stereo Event-based Particle Tracking Velocimetry for 3D Fluid Flow Reconstruction

Yuanhao Wang, Ramzi Idoughi, and Wolfgang Heidrich

King Abdullah University of Science and Technology, Thuwal, Saudi Arabia
{yuanhao.wang, ramzi.idoughi, wolfgang.heidrich}@kaust.edu.sa

Abstract. Existing Particle Imaging Velocimetry techniques require the use of high-speed cameras to reconstruct time-resolved fluid flows. These cameras provide high-resolution images at high frame rates, which generates bandwidth and memory issues. By capturing only changes in the brightness with a very low latency and at low data rate, event-based cameras have the ability to tackle such issues. In this paper, we present a new framework that retrieves dense 3D measurements of the fluid velocity field using a pair of event-based cameras. First, we track particles inside the two event sequences in order to estimate their 2D velocity in the two sequences of images. A stereo-matching step is then performed to retrieve their 3D positions. These intermediate outputs are incorporated into an optimization framework that also includes physically plausible regularizers, in order to retrieve the 3D velocity field. Extensive experiments on both simulated and real data demonstrate the efficacy of our approach.

Keywords: Fluid imaging, Event-based camera, Particle imaging velocimetry, Stereo-PTV, Optimization

1 Introduction

Fluid imaging is a topic of interest for several scientific and engineering areas like fluid dynamic, combustion, biology, computer vision and graphics. The capture of the 3D fluid flow is a common requirement to characterize the fluid and its motion regardless the application domain. Despite the large number of contributions to this field, retrieving a 3D dense measurement of the velocity vector over the fluid remains a challenging task. Different techniques have been proposed to capture and measure the fluid motion. The most commonly used approach involves introducing tracers such as dye, smoke, or particles into the studied fluid. Then by tracking the advected motion of the tracers, the fluid flows are retrieved. Particle Tracking Velocimetry (PTV) and Particle Imaging Velocimetry (PIV) are the two most popular techniques among the tracer-based approaches [1,52,13]. The PTV methods follow a Lagrangian formalism, where each particle is tracked individually. On the other hand, the PIV techniques retrieve the velocity in an Eulerian fashion by tracking the particles as a group, like a texture patch.

In basic PIV approaches [46], a thin slice of the fluid is illuminated by a laser sheet. From an in-plane tracking of the illuminated particles, a dense measurement of two components of the motion field can be retrieved inside the illuminated slice of the volume. However, turbulent or unsteady flows cannot be fully characterized with such planar measurements. To solve this issue, several variants of PIV have been proposed to extend the velocity retrieval to the third spatial dimension (a detailed discussion about such techniques is provided in the next section). Overall, these techniques can be regrouped into two families: multiple-cameras based approaches like tomographic PIV (tomo-PIV) [18,61] and single-camera based techniques such as light-field (or plenoptic) PIV [19,66,42] or structured-light PIV [70,69,2]. Among these, tomo-PIV is considered the established reference technology for the velocity 3D measurement, since it provides high spatial and temporal resolutions. However, it requires to have a precise calibration and synchronization of the used cameras. Moreover, for the tomo-PIV setups, the depth-of-field is a real limitation on the size of the volume of interest. On top of that, setup needed to reach a high temporal resolution can be very costly. Finally, due to space limitation on the hardware setup, only a few cameras (4-6) can be used, which limits the reconstruction quality. Although the proposed single-camera based techniques overcame the main shortcomings of the tomo-PIV (calibration, synchronization and the space limitation), they still have some limitations. The plenoptic PIV systems have significantly lower spatial resolution. Moreover, the current light-field cameras have low frame rates, which reduces the temporal resolution of the retrieved flow field. Furthermore, the main limitation of the structured-light PIV methods, like the RainbowPIV [70], is also the limited spatial resolution along the axial dimension.

To deal with many common flows, the cameras used in PIV systems should have a sufficiently high frame rate. However, the high speed cameras are relatively expensive, which particularly impacts the cost of multi-camera setups like tomo-PIV. The large bandwidth and storage requirements of such solutions pose additional difficulties. In order to address these issues, this paper introduces a new stereo-PIV technique based on two event cameras. Thanks to some outstanding properties like a very low latency ($1 \mu s$) and a low power consumption [30,49,8,64], event cameras are well suited for fast motions detection and tracking. After capturing two sequences of event images using two different cameras, we propose a new framework that retrieves the velocity field of the fluid with a high time resolution. Indeed, the events based cameras can react very fast to motions given their very low latency. Thus, no motion blur can be observed with such cameras. In addition, we have a complete control over the time interval over which the events are aggregated.

The main technical contributions of this work are:

1. To the best of our knowledge, we propose the first event-camera based stereo-PIV setup for measuring time-resolved fluid flows.
2. We formulate a pertinent data term that links the event images to the 3D fluid velocity vector field.

3. We propose an optimization framework to retrieve the fluid velocity field from the event images. This framework include physically-based priors to solve the ill-posed inverse problem.
4. We demonstrate the accuracy of our approach on both simulated and real fluid flows.

2 Related work

Fluid imaging is an active and challenging research topic for several domains such as fluid dynamics, combustion, biology, computer vision and graphics. To capture a fluid and its flow, several techniques have been applied in order to retrieve some of the fluid characteristics like the temperature, the density, the species concentration (scalar fields), the velocity or the vorticity (vector fields) [67]. In computer vision and graphics, an initial effort was focused on retrieving some physical properties of the fluid, from which a good visualization can be obtained. Thus, the captured properties vary from one fluid to another. As examples, light emission, refractive index, scattering density, and dye concentration have been reconstructed to visualize flames [33,26], air plumes [34,4], liquid surfaces [32,68,44], smoke [27,25] and fluid mixtures [24,23]. More recently, the interest has shifted to the estimation of velocity field, in order to improve the scalar density reconstruction [68,16,17,72], or as the final output [4,23,70,69,37].

Fluid velocity estimation techniques are mostly tracer-based approaches. Tracers like particles or dye are introduced into the fluid. Then, the velocity field of the fluid is recovered by tracking those tracers. However, tracer-free methods have been also investigated. Examples of such methods, Background Oriented Schlieren (BOS) [56,4,51] and Schlieren PIV [35,10], use the variations of the refractive index of the fluid as a "tracer" of the fluid motion.

Among the tracer-based methods, Particle Imaging Velocimetry (PIV) is the most commonly used [1,52]. During the last three decades, several techniques have been proposed to extend the standard PIV [46], where only two components of the velocity are measured in a thin slice of the volume (2D-2C). Stereoscopic PIV (Stereo-PIV) [50] records the region of interest using two synchronized cameras, which allows to retrieve the out-of-plane velocity component (2D-3C). The 3D scanning PIV (SPIV) technique [12,31], performs a standard PIV reconstruction on a large set of parallel light-sheet planes that samples the volume. A fast scanning laser is used to illuminate those planes at a high scanning rate. This approach however still retrieves only two-dimensions of the velocity (3D-2C). Moreover, at each time only one depth layer is scanned. More sophisticated techniques that can resolve 3D volumes include holographic PIV [43,29], defocusing digital PIV [47,48,71], synthetic aperture PIV [6], tomographic PIV (tomo-PIV) [18,61], structured-light PIV [70,69,2], and plenoptic (light-field) PIV [19,66,42]. These approaches can be multi-view based, like the widely used tomo-PIV. In this case, the hardware setup induces some difficulties like the calibration and synchronization of the cameras, as well as a limited

space. Moreover, the reconstructed volume is usually small, since it should be included in the field of view of all cameras. On the other hand, mono-camera PIV approaches encode the depth information using color/intensity or using light-field camera. In the case of Rainbow PIV [70], the lower sensitivity of the camera to the wavelength change and the light scattering limit the depth resolution of the retrieved velocity field. Similar drawback can be observed with the intensity-coded PIV setup [2]. Otherwise, the plenoptic PIV systems have a more limited spatial resolution, since they capture on the same image several copies of the volume from different angles. They also suffer from a reduced temporal resolution, given the frame rate of current light-field cameras.

On the other hand, Particle Tracking Velocimetry (PTV) techniques [13,39] work by tracking individual particles to retrieve their velocity. The obtained flow field is then sampled according to these particles. Some variational approaches [28,21,63] were proposed to optimize the 3D flow field from particle tracks in order to estimate the flow in all the volume. Another advantage of using variational approaches is to incorporate physical constraints as prior in the optimization framework. This has been done in several PTV and PIV techniques [58,60,59,3,28,70,38].

The major limitation of PTV approaches has been the low number of particles that can be tracked. However, the recent shake-the-box method [62] succeeded into tracking particles with high densities, similar to those in PIV techniques.

All of these techniques require the use of high speed cameras, in order to reconstruct many real-world flow phenomena. Combined with the high resolution requirement, the large data generated at each capture (typically a rate of 2 GB/s) induces high bandwidth and large memory specifications for the used camera(s). To address this issue, [15] presented a proof-of-concept study about the use of dynamic vision sensor (DVS) in the capture and the tracking of particles. The proposed algorithm is suitable only for the 2D tracking of a sparse set of fully resolved particles (10 pixels diameter) in the volume. Recently, [11] proposed another approach based on Kalman filters to track neutrally buoyant soap bubbles from three cameras. This technique can handle the tracking and 3D reconstruction of under-resolved particles, which allows the larger studied volumes. However, this approach is not well-suited for high particle seeding densities, and it is not able to provide a dense measurement of the velocity field. In this paper, we propose a new framework that reconstructs a dense time-resolved 3D fluid flows from two event-based cameras.

Event cameras (a.k.a dynamic vision sensors) were first developed by [41] to mimic the retina of eyes, which is more sensitive to motions. These cameras respond only to the brightness change in the scene asynchronously and independently for each pixel. Event cameras have a very low latency (up to 1 μs), a low power consumption, and a high dynamic range [30,49,8,64]. These properties are major assets to fulfill several computer vision tasks. For instance, event cameras were introduced for object and feature tracking [14,22], depth estimation [57,53], optical flow estimation [7,5,74], high dynamic range imaging [55] and many other applications. An exhaustive survey about event camera appli-

cations can be found in [20]. In our framework, the event cameras are used for the particle tracking and optical flow estimation.

To track micro-particles using an event-camera, [45] estimate the particles' positions using an event-based Hough circle transform, combined with a centroid (a centre-of-mass algorithm). [9] apply an event-based visual flow algorithm [7] in order to track particles imaged by a full-field Optical Coherence Tomography (OCT) setup. This visual flow algorithm estimates the normal flow by fitting the events to a plane in the x-y-t space. The optical flow is then estimated as the slope of this plane. However, these two approaches reconstruct only 2D velocities, and are limited to the tracking of sparse particle densities.

To improve the optical flow estimation from event sequences, [73,75] propose to regroup events into features in a probabilistic way. This assignment is governed by the length of the optical flow. The latter is computed as a maximization of the expectation of all these assignments. An affine fit is used to model the features deformation. In our approach we use this technique to compute the optical flow over the event sequences.

3 Proposed method

3.1 System overview

We propose a new particle tracking velocimetry technique for the reconstruction of dense 3D fluid flows captured by two event-cameras. Our framework, illustrated in Figure 1, is mainly composed of four modules: (1) a **Cameras calibration step**, which entails estimation of the camera calibration matrices for both event cameras. (2) **Event feature tracking** for a 2D particle velocity reconstruction. In this step, we apply feature tracking algorithms proposed by [73,75] in order to track the particles in the two captured sequences of event images. Then the 2D particle velocity in the two cameras image planes is recovered. (3) A **Stereo matching step** is performed using a double triangulation method to find the position of the particles in the 3D volume. (4) **3D velocity field reconstruction**. This last module is an optimization framework that includes our derived data fitting term, and a physically constrained 3D optical flow model. In the following section, we provide a detailed presentation of the three main modules of our framework.

3.2 Image formation model

In this section, we derive a model that links the 3D fluid velocity field to the captured event sequences.

Event camera model. Each pixel (x, y) of an event camera independently generates an event $e_i = (x, y, t_i, \rho_i)$, when detecting a brightness change higher than a pre-defined threshold τ :

$$|L(x, y, t_i) - L(x, y, t_{previous})| \geq \tau \quad (1)$$

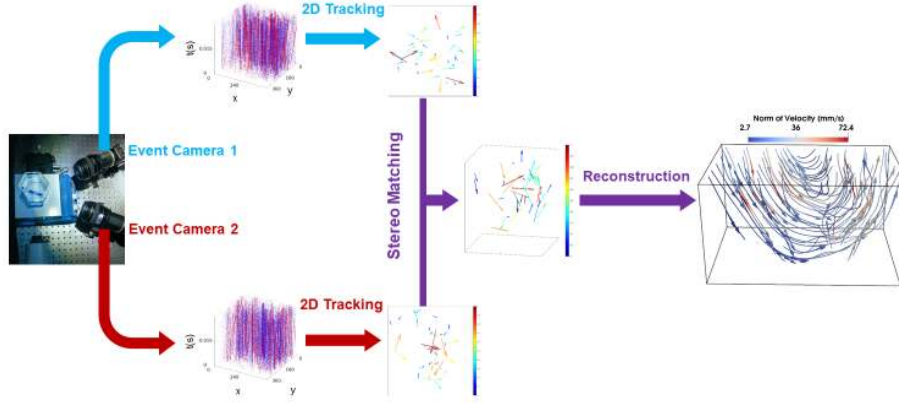


Fig. 1: Overview of the architecture of our stereo-event PTV framework. The two event cameras capture the motion of the particles inside the fluid. They generate two sequence of events, represented here in the x - y - t space. A 2D tracking step provide the 2D velocity of the captured particles for each sequence. Then, using a stereo matching step we build a sparse 3D velocity field that we use in order to estimate the dense 3D fluid flow.

where $L(x, y, t_i)$ is the brightness (log intensity) of the pixel (x, y) at time t_i , $t_{previous}$ is the time of occurrence of the previous event and $\rho_i = \pm 1$ is the event polarity corresponding to the sign of the brightness change.

Camera calibration. Each point (X, Y, Z) in the fluid is projected onto a pixel (x^k, y^k) of the k^{th} camera image plane:

$$\alpha \cdot [x^k, y^k, 1]^T = \mathbf{M}_k [X, Y, Z, 1]^T \quad (2)$$

where α is a scale factor, and \mathbf{M}_k is a 4×3 matrix describing the k^{th} camera calibration matrix (intrinsic + extrinsic parameters). This matrix is obtained during the calibration step. Note that we use lowercase letters for 2D pixel coordinates and the Uppercase for the 3D voxel positions.

These camera calibration matrices are also used to project the 3D fluid velocity field $\mathbf{u} = [\mathbf{u}_X, \mathbf{u}_Y, \mathbf{u}_Z]^T$ onto the image planes of the cameras. The obtained 2D velocity fields are denoted \mathbf{v}_k . Given that the velocity can only be measured for the positions where particles are present, we can write:

$$\mathbf{v}_k = \mathbf{p}_k \odot (\mathbf{M}_k \mathbf{u}) \quad (3)$$

where \odot is the Hadamard product, and \mathbf{p}_k is the particle occupancy distribution. \mathbf{p}_k equals to 1 when a particle is mapped to the k^{th} camera.

The 3D velocity field can be retrieved by solving the linear system given in Equation 3 for the two event cameras. However, the projected velocities \mathbf{v}_k

are not directly obtained from the event cameras. In the following section, we explain the approach used to estimate these velocities from the event sequences.

3.3 Event-based particle tracking

The main objective of this step is to recover (\mathbf{v}_1 and \mathbf{v}_2), the 2D velocity fields of the particles in the two captured event sequences. First, we pre-process the event sequences by applying a circular averaging filter in order to simplify the detection of the particles centers. The size of this filter is chosen according to the particle size in the images. At this stage the center coordinates of each particle in the two event sequences can be easily determined. After this preprocessing step, we use the event-based optical flow method introduced by [73,75], in order to track the particles in the image planes, and then retrieve the velocities \mathbf{v}_1 and \mathbf{v}_2 . In this approach, the events $\{e_i\}$ are associated with a set of features representing the particles in our case. All events associated with a given particle P are within a spatio-temporal window, where the average flow $v(P, t)$ is assumed to be constant if the temporal dimension $[t, t + \Delta t]$ of the window is small enough. This window can be written as:

$$W(R, t) := \{e_i \mid t_i \leq t, \|(x_i, y_i) - t_i v(P, t) - (x_c, y_c)\| \leq R\}, \quad (4)$$

where R and (x_c, y_c) are respectively the spatial extension (radius) and the coordinate of the particle center in the image plane.

The association of events to particles is defined in terms of the flow $v(P, t)$ that we would like to estimate: events corresponding to the same 3D point should propagate backward onto the same image position. [73] propose an Expectation Maximization algorithm to solve this flow constraint. In the first step (E step), they update the association between events and the particles, given a fixed flow $v(P, t)$. Then, this flow is updated in the second step (M step), using the new matches between events and particles. More details about the implementation of this algorithm can be found in [73,75]. By applying this algorithm for tracking all particles and for all time stamps, we can recover the 2D velocity fields \mathbf{v}_1 and \mathbf{v}_2 .

3.4 Stereo matching

The aim of this step is on one hand to find the particle positions in 3D space. On the other hand, the retrieved 2D velocity field can be backprojected to get an estimation of the 3D velocity field \mathbf{u} for the positions corresponding to the particles. We perform this stereo matching step using a triangulation procedure [36,40]. The main idea is to build for each particle a pixel-to-line transformation and then find the 3D positions that minimizes the total distances to all the lines.

For each identified particle P_i in an image captured by camera 1, its center $(x_{i,c}^1, y_{i,c}^1)$ is used to backproject a line of sight to the different planes of the volume of interest. Then, the intersection points of this line with the different

planes of the volume will be reprojected on the corresponding image frame captured by camera 2. Candidate particles P_j are selected in camera 2 only if the distance d_{ij}^{1-2} between the particle’s center on the camera 2 and the reprojected points corresponding to the center of particles P_i on the camera 1 are under a given threshold distance (2 pixels, for example). Similarly, we perform the inverse mapping, a particle P_i on the camera 2 is backprojected to the fluid volume then reprojected to the image plane of the camera 1 to find candidates particles P_i , under the constraint that the distance d_{ij}^{2-1} is less than the fixed threshold. The correspondence between the two cameras are obtained by minimize the summation of all the distances. This is formulated as a simple linear assignment problem, and solved by using the Hungarian algorithm:

$$\begin{aligned} & \min C_{ij}d_{ij}^{1-2} + C_{ij}d_{ij}^{2-1} & (5) \\ & \text{subject to } \begin{cases} \sum_j C_{ij} \leq 1 \\ \sum_i C_{ij} \leq 1 \\ C_{ij} \in \{0, 1\} \end{cases} \end{aligned}$$

From this stereo matching step, we estimate the particles 3D position as well as the velocity of those particles. However in practice, because of occlusions and noise, some of the particles may not be matched or, worse, they might be mismatched. This is what motivate us to use a variational approach to improve the particles’ velocity estimation and also to extend the velocity estimation to whole volume of interest.

3.5 3D velocimetry reconstruction

We propose to reconstruct the 3D fluid velocity field $\mathbf{u} = [\mathbf{u}_X, \mathbf{u}_Y, \mathbf{u}_Z]^T$ for each voxel of the volume of interest, by solving Equation 3 using the two event cameras, and by combining all time frames. In order to handle this ill-posed inverse problem, we introduce several regularizer terms, directly derived from the physical properties of the fluid.

Data fitting-term. As mentioned previously, we define the data-fitting term from Equation 3. This term translates that the projection of the 3D velocity field to each camera image plane, should be consistent with the 2D velocity field observed with that camera. The data-fitting term can then be written as follows:

$$E_{data}(\mathbf{u}) = \frac{1}{2}\|\mathbf{p} \odot (\mathbf{M}_1\mathbf{u}) - \mathbf{v}_1\|_2^2 + \frac{1}{2}\|\mathbf{p} \odot (\mathbf{M}_2\mathbf{u}) - \mathbf{v}_2\|_2^2, \quad (6)$$

where $\mathbf{p} = \mathbf{p}_1 \odot \mathbf{p}_2$ is the particle occupancy distribution, that take into account only the matched particles between the 2 cameras.

Spatial smoothness. The second term of our optimization is a spatial smoothness on the 3D velocity field. The advantage of this term is to help in giving a better interpolation for the voxels where no particles are detected.

$$E_{smooth}(\mathbf{u}) = \|\nabla_S \mathbf{u}\|_2^2 \quad (7)$$

Incompressibility. In the case of incompressible fluid simulation or capture, it is common to constrain the flow field to be divergence-free [23,70,16]. This constraint is derived directly from the mass-conservation equation for the fluid. Usually the divergence-free regularization is applied by projecting the velocity field onto the space of divergence-free velocity field. However, we notice that in the case of lower spatial resolution, the discretization of the divergence operator may introduce some divergence to the flow. Therefore, we prefer to include the incompressibility prior as a soft (L-2) constraint instead of a hard projection.

$$E_{div}(\mathbf{u}) = \|\text{div}(\mathbf{u})\|_2^2 \quad (8)$$

Temporal coherence. In the absence of external forces, the Navier-Stokes equation for a non-viscous fluid can be simplified as follows:

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} = 0 \quad (9)$$

This equation can be used as an approximation of the temporal evolution of the fluid flow. We can then, advect the velocity field at a given time stamp by itself to deduce an estimate of the field at the next time stamp. This advection is applied in a forward and backward manners. This yields the following term:

$$E_{TC}(\mathbf{u}_t) = \|\mathbf{p} \odot (\mathbf{u}_t - \text{advect}(\mathbf{u}_{t-1}, \Delta t))\|_2^2 + \|\mathbf{p} \odot (\mathbf{u}_t - \text{advect}(\mathbf{u}_{t+1}, -\Delta t))\|_2^2 \quad (10)$$

where \mathbf{u}_t is the velocity field at the time t . The particle occupancy \mathbf{p} is used here as a mask, in order to take into account only the regions of the volume where particles have been observed, since these are the only regions with reliable velocity estimates.

Optimization framework. The general optimization framework is then expressed as:

$$(\mathbf{u}^*) = \underset{\mathbf{u}}{\text{argmin}} E_{data}(\mathbf{u}) + \lambda_1 E_{smooth}(\mathbf{u}) + \lambda_2 E_{TC}(\mathbf{u}) + \lambda_3 E_{div}(\mathbf{u}) \quad (11)$$

Equation 11 is composed only of L-2 terms. We solve it then using the conjugate gradient method. To handle the large velocities in the fluid, we build a multi-scale coarse-to-fine scheme [70]. More details about our implementation are given in the supplement material.

4 Experiments

Experiments on both simulated and captured fluid flows were conducted to evaluate our approach. We implemented our framework in matlab ¹. All the experiments were conducted on a computer with an Intel Xeon E5-2680 CPU processor and 128 GB RAM. The reconstruction time for the simulated dataset which contains 8 frames (with $80 \times 80 \times 80$ voxels) was around 28 minutes. Furthermore, the parameter settings for the optimization (see Equation 11) were kept the same for all the datasets ($\lambda_1 = 2.5 \times 10^{-5}$, $\lambda_2 = 0.025$, $\lambda_3 = 2.5 \times 10^{-5}$).

4.1 Synthetic data

To quantitatively assess our method, we simulated a fluid undergoing a rigid-body-like vortex, with a fixed angular speed. A volume with a size of $20 \text{ mm} \times 20 \text{ mm} \times 20 \text{ mm}$ was seeded randomly with particles of an averaged size of 0.1 mm (1 % variance). This fluid is captured by two different simulated event cameras having a spatial resolution of 800×800 , which is similar to the real experiment setup in Figure 5. Different vortex speeds and particle densities were simulated. The approach introduced by [65] was applied to advect the particles over time using the vortex velocity field. Moreover, we simulated different frame rate images. Finally, we used E-sim code [54] to generate the event sequences observed by the simulated sensors.

Ablation study. In order to illustrate the impact of each of our priors, we conducted an ablation study. We compare our method without the use of the temporal coherence and the divergence terms (**w/o** \mathbf{E}_{TC} & \mathbf{E}_{div}), our method without the divergence term (**w/o** \mathbf{E}_{div}), and our proposed method (**Ours**). For the quantitative comparison, we use two metrics: the average angular error (AAE), i.e. the average discrepancy in the flow direction, and the average end-point error (AEE), i.e. the average Euclidean norm of the difference between the real and estimated flow vectors.

In Figure 2, we illustrate the velocity field reconstruction using our method versus the ground truth. Except for the borders, the reconstruction is very accurate. The numerical results of the ablation study are shown in the Figure 3. As expected, both the AAE and the EPE errors are improved when adding the different priors. Moreover, these errors are almost constant from one time frame to another. We need to point out that the temporal coherence term might not improve too much the reconstruction for all frames. However, in the general case, it smoothes the result in the temporal domain, which is important for visual quality in frame-based or time-based data processing.

In Figure 4 we illustrate for a 2D slice the end point error as well as the divergence of the velocity field for different methods. The mean error for the different methods is 0.182, 0.178, 0.171 respectively. The error will generally become smaller gradually as expected. The mean absolute divergence for the three different method is 0.0096, 0.0100, 0.0071. We notice that the temporal

¹ The code is available on: <https://github.com/vccimaging/StereoEventPTV>.

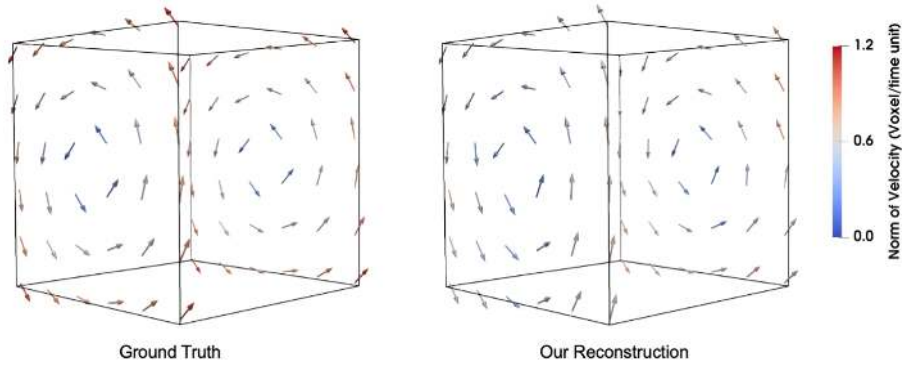


Fig. 2: Ground truth (left) and reconstruction result using our method (right) for a simulated rigid-body-like vortex.




Fig. 3: Quantitative comparisons with ground truth velocity field for different reconstruction using different priors. Left: Average angular error (in degree). Right: Average end-point error (in voxel).




Fig. 4: End point error (first row) and the divergence of the flow (second row) computed on a 2D slice. From left to right: Ground truth, Our method without the incompressibility and the temporal coherence terms, Our method without the incompressibility term, and Our proposed method

coherence term introduces some divergence to the flow. It can be explained by the fact that the temporal smoothness might propagate wrong stereo matching

to adjacent time frame. However, the incompressibility constraint will reduce the divergence and bring it closer to zero.

Table 1: Quantitative evaluation (AAE in degree / EPE in voxel) of different particle densities and at different rotation speeds.

Angular Speed (rad/s)	Particle density (ppp)				
	0.006	0.012	0.018	0.024	0.03
5	9.73/0.046	10.22/0.042	9.56/0.043	15.38/0.052	12.83/0.048
10	10.20/0.097	11.12/0.090	10.72/0.092	13.34/0.097	12.12/0.087
15	10.58/0.146	10.38/0.131	9.94/0.123	13.77/0.150	11.77/0.133
20	9.73/0.187	10.37/0.164	9.76/0.168	12.13/0.178	11.49/0.180

Particle densities and vortex speed impact. We also evaluated our method for different particle densities and different angular speeds of the vortex. The results are shown in Table 1. These experiments have been conducted for the same duration. As expected the larger the speed, the larger the EPE. However, the angular error is in the same range independent of the vortex speed. On the other hand, from these experiments we can deduce that our method can handle a wide range of particle densities. These experiments show that our method can be used in very different situations with a wide range of particle densities and different fluid velocities.

4.2 Captured data

Experimental setup The experimental setup used for the event-based fluid imaging is shown in Figure 5. To capture the stereo events at the same time, we utilized two synchronized Prophesee cameras (Model: PEK3SHEM, Sensor: CSD3SVCD [$7.2\text{ mm} \times 5.4\text{ mm}$], 480×360 pixels), with an angle of 60 degrees between the two optical axes. Two lenses with a focal length of 85 mm and a 3D printed extension tubes were attached to the cameras. The aperture was set to $f/16$ to have a depth-of-field of 10 mm . The tank was seeded with white particles (White Polyethylene Microspheres) having a diameter in the range $[90, 106\ \mu\text{m}]$. The size of the particles on the image plane is approximately 6.7 pixels. By applying downsampling with a downsample factor of 6 and stereo matching, we reconstruct a volume with: $78 \times 48 \times 42$ voxels. For the calibration step, a 17×16 checkerboard where each square has an edge length of 0.5 mm was attached on a glass slide. We used a controllable translation stage to modify the distance of the checkerboard to the cameras. More details about the calibration can be found in the supplement material.

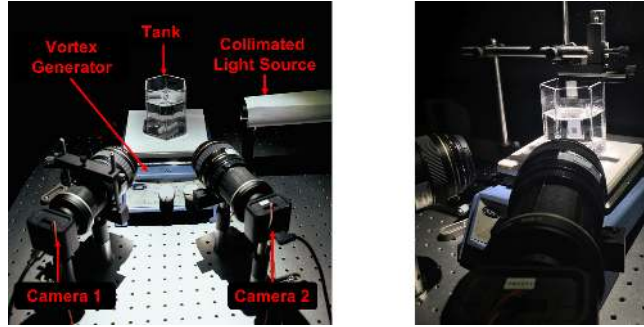


Fig. 5: Left: Illustration of our experimental setup. A collimated white light source illuminate the hexagonal tank. A vortex generator is used to control the speed of the vortex during the experiments. Right: illustration of the calibration step, where images of a small check board are captured for several positions. A controlled translation stage is used to change the positions.

Controlled vortex flow The first experiment we performed was a controlled vortex flow. We used a magnetic stirring rod (Model: Stuart CB162) to generate different vortices by controlling the rotation speed of the stirring rod. We evaluate our reconstruction method over the different vortices. The reconstructed streamlines for two examples are shown in Figure 6. We can see that our reconstruction offers a good representation of the vortex structure, and the velocity norm seems to be reliable given the speed of the stirring rod. Please refer to the supplement for more results.

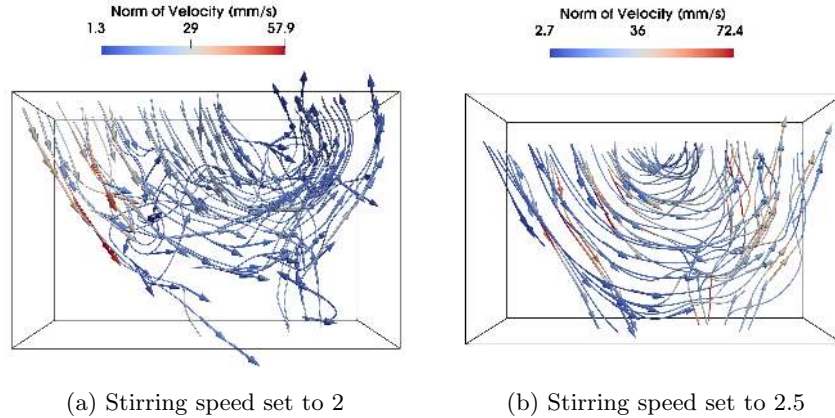
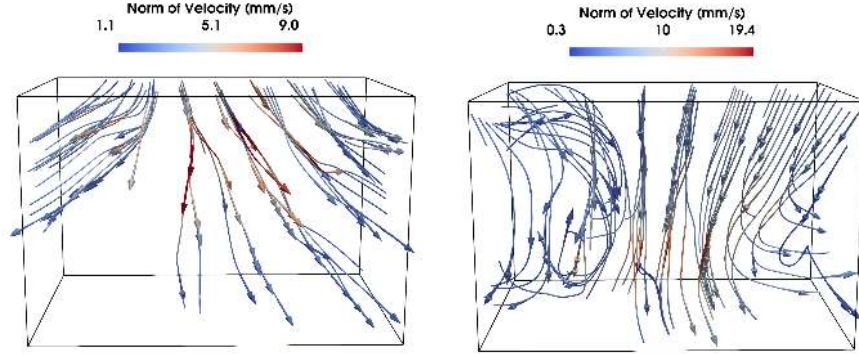


Fig. 6: Streamline visualization for controlled vortex flows. Left: The stirrer speed was set to 2. Right: The stirrer speed was set to 2.5.

Fluid injection Finally, we conducted another set of experiments, consisting of a relatively fast fluid injections into the tank using a syringe. As shown in Figure

7, different speeds of the flow and different injection directions can be easily distinguished from our reconstructed results. Additional results and illustrations are presented in the supplemental material.



(a) Deflected injection.

(b) Vertical injection.

Fig. 7: Streamline visualization for a fast fluid injections using a syringe. By controlling the syringe orientation we have captured: (a) a deflected injection and (b) a vertical injection.

5 Conclusions

We have introduced a stereo event-based camera system coupled with 3D fluid flow reconstruction strategies in this paper. Instead of using image based optical flow reconstruction in the traditional tomographic PTV, our approach is based on generating the two dimensional flow from the event information, and then matching the resulting trajectories in 3D to obtain full 3D-3C flow fields.

Both the numerical and experimental assessment confirm the effectiveness of our approach. By simulating different particle numbers in the tank that usually used in the PTV, we found that our method works on a wide range of particle densities. Furthermore, by controlling the stirring speed of the vortex, we found that our approach can deal with fast fluid flow.

There are some drawbacks to our approach. First of all, the spatial resolution of currently available event cameras is quite low, which also adversely impacts the spatial resolution of the reconstruction. Second, due to the high dynamic range of the event camera, the light intensity and the camera sensitivity should be carefully selected to have a good measurements in real experiments. Last but not least, the bandwidth of the event-camera is limited, the method fails when the speed of the controlled vortex exceeds a certain threshold, in which case the bus of the camera was saturated. However, with future improvements of event camera hardware, we believe these shortcomings can be overcome, making our method an attractive option for 3D-3C fluid imaging.

References

1. Adrian, R.J., Westerweel, J.: Particle image velocimetry. Cambridge University Press (2011)
2. Aguirre-Pablo, A., Aljedaani, A.B., Xiong, J., Idoughi, R., Heidrich, W., Thoroddsen, S.T.: Single-camera 3D PTV using particle intensities and structured light. *Experiments in Fluids* **60**(2), 25 (2019)
3. Álvarez, L., Castano, C., García, M., Krissian, K., Mazorra, L., Salgado, A., Sánchez, J.: A new energy-based method for 3D motion estimation of incompressible PIV flows. *Computer Vision & Image Understanding* **113**(7), 802–810 (2009)
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. *ACM Trans. Graph.* **27**(5), 132 (2008)
5. Bardow, P., Davison, A.J., Leutenegger, S.: Simultaneous optical flow and intensity estimation from an event camera. In: *Proc. CVPR*. pp. 884–892 (2016)
6. 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)
7. Benosman, R., Clercq, C., Lagorce, X., Ieng, S.H., Bartolozzi, C.: Event-based visual flow. *IEEE transactions on neural networks and learning systems* **25**(2), 407–417 (2013)
8. Berner, R., Brandli, C., Yang, M., Liu, S.C., Delbruck, T.: A 240×180 10mW $12\mu\text{s}$ latency sparse-output vision sensor for mobile applications. In: *2013 Symposium on VLSI Circuits*. pp. C186–C187. IEEE (2013)
9. Berthelon, X., Chenegros, G., Libert, N., Sahel, J.A., Grieve, K., Benosman, R.: Full-field OCT technique for high speed event-based optical flow and particle tracking. *Optics express* **25**(11), 12611–12621 (2017)
10. Biswas, S.: Schlieren image velocimetry (SIV). In: *Physics of Turbulent Jet Ignition*, pp. 35–64. Springer (2018)
11. Borer, D., Delbruck, T., Rösgen, T.: Three-dimensional particle tracking velocimetry using dynamic vision sensors. *Experiments in Fluids* **58**(12), 165 (2017)
12. 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)
13. Dabiri, D., Pecora, C.: Particle Tracking Velocimetry. IOP Publishing (2020)
14. Delbruck, T., Lang, M.: Robotic goalie with 3 ms reaction time at 4% CPU load using event-based dynamic vision sensor. *Frontiers in neuroscience* **7**, 223 (2013)
15. Drazen, D., Lichtsteiner, P., Häfziger, P., Delbrück, T., Jensen, A.: Toward real-time particle tracking using an event-based dynamic vision sensor. *Experiments in Fluids* **51**(5), 1465 (2011)
16. Eckert, M.L., Heidrich, W., Thürey, N.: Coupled fluid density and motion from single views. In: *CGF*. vol. 37, pp. 47–58. Wiley Online Library (2018)
17. Eckert, M.L., Um, K., Thürey, N.: ScalarFlow: a large-scale volumetric data set of real-world scalar transport flows for computer animation and machine learning. *ACM Trans. Graph.* **38**(6), 1–16 (2019)
18. Elsinga, G.E., Scarano, F., Wieneke, B., van Oudheusden, B.W.: Tomographic particle image velocimetry. *Experiments in Fluids* **41**(6), 933–947 (2006)
19. Fahringer, T.W., Lynch, K.P., Thurow, B.S.: Volumetric particle image velocimetry with a single plenoptic camera. *Measurement Science and Technology* **26**(11), 115201 (2015)

20. Gallego, G., Delbruck, T., Orchard, G., Bartolozzi, C., Taba, B., Censi, A., Leutenegger, S., Davison, A., Conradt, J., Daniilidis, K., et al.: Event-based vision: A survey. arXiv preprint arXiv:1904.08405 (2019)
21. Gesemann, S., Huhn, F., Schanz, D., Schröder, A.: From noisy particle tracks to velocity, acceleration and pressure fields using B-splines and penalties. In: 18th international symposium on applications of laser and imaging techniques to fluid mechanics, Lisbon, Portugal. pp. 4–7 (2016)
22. Glover, A., Bartolozzi, C.: Event-driven ball detection and gaze fixation in clutter. In: Proc. IROS. pp. 2203–2208. IEEE (2016)
23. Gregson, J., Ihrke, I., Thuerey, N., Heidrich, W.: From capture to simulation: connecting forward and inverse problems in fluids. *ACM Trans. Graph.* **33**(4), 139 (2014)
24. Gregson, J., Krimerman, M., Hullin, M.B., Heidrich, W.: Stochastic tomography and its applications in 3D imaging of mixing fluids. *ACM Trans. Graph.* **31**(4), 52–1 (2012)
25. Gu, J., Nayar, S.K., Grinspun, E., Belhumeur, P.N., Ramamoorthi, R.: Compressive structured light for recovering inhomogeneous participating media. *IEEE Trans. PAMI* **35**(3), 1–1 (2012)
26. Hasinoff, S.W., Kutulakos, K.N.: Photo-consistent reconstruction of semitransparent scenes by density-sheet decomposition. *IEEE Trans. PAMI* **29**(5), 870–885 (2007)
27. Hawkins, T., Einarsson, P., Debevec, P.: Acquisition of time-varying participating media. Tech. rep., UNIVERSITY OF SOUTHERN CALIFORNIA MARINA DEL REY CA INST FOR CREATIVE (2005)
28. Heitz, D., Mémin, E., Schnörr, C.: Variational fluid flow measurements from image sequences: synopsis and perspectives. *Experiments in fluids* **48**(3), 369–393 (2010)
29. Hinsch, K.D.: Holographic particle image velocimetry. *Measurement Science and Technology* **13**(7), R61 (2002)
30. Hofstatter, M., Schön, P., Posch, C.: A SPARC-compatible general purpose address-event processor with 20-bit 10ns-resolution asynchronous sensor data interface in 0.18 μm CMOS. In: Proceedings of 2010 IEEE International Symposium on Circuits and Systems. pp. 4229–4232. IEEE (2010)
31. Hori, T., Sakakibara, J.: High-speed scanning stereoscopic PIV for 3D vorticity measurement in liquids. *Measurement Science and Technology* **15**(6), 1067 (2004)
32. Ihrke, I., Goidluecke, B., Magnor, M.: Reconstructing the geometry of flowing water. In: Proc. ICCV. vol. 2, pp. 1055–1060. IEEE (2005)
33. Ihrke, I., Magnor, M.: Image-based tomographic reconstruction of flames. In: Proc. SCA. pp. 365–373 (2004)
34. Ji, Y., Ye, J., Yu, J.: Reconstructing gas flows using light-path approximation. In: Proc. CVPR. pp. 2507–2514 (2013)
35. Jonassen, D.R., Settles, G.S., Tronosky, M.D.: Schlieren piv for turbulent flows. *Optics and Lasers in Engineering* **44**(3-4), 190–207 (2006)
36. Knutsen, A.N., Lawson, J.M., Dawson, J.R., Worth, N.A.: A laser sheet self-calibration method for scanning PIV. *Experiments in Fluids* **58**(10), 145 (2017)
37. Lasinger, K., Vogel, C., Pock, T., Schindler, K.: 3D fluid flow estimation with integrated particle reconstruction. *Int. J. Computer Vision* pp. 1–16 (2019)
38. Lasinger, K., Vogel, C., Schindler, K.: Volumetric flow estimation for incompressible fluids using the stationary stokes equations. In: Proc. ICCV. pp. 2565–2573 (2017)
39. Maas, H., Gruen, A., Papantoniou, D.: Particle tracking velocimetry in three-dimensional flows. *Experiments in fluids* **15**(2), 133–146 (1993)

40. Machicoane, N., Aliseda, A., Volk, R., Bourgoïn, M.: A simplified and versatile calibration method for multi-camera optical systems in 3D particle imaging. *Review of Scientific Instruments* **90**(3), 035112 (2019)
41. Mahowald, M.: VLSI analogs of neuronal visual processing: a synthesis of form and function. Ph.D. thesis, California Institute of Technology Pasadena (1992)
42. Mei, D., Ding, J., Shi, S., New, T., Soria, J.: High resolution volumetric dual-camera light-field PIV. *Experiments in Fluids* **60**(8), 132 (2019)
43. Meng, H., Hussain, F.: Holographic particle velocimetry: a 3D measurement technique for vortex interactions, coherent structures and turbulence. *Fluid Dynamics Research* **8**(1-4), 33 (1991)
44. Morris, N.J., Kutulakos, K.N.: Dynamic refraction stereo. *IEEE transactions on pattern analysis and machine intelligence* **33**(8), 1518–1531 (2011)
45. Ni, Z., Pacoret, C., Benosman, R., Ieng, S., RÉGNIER*, S.: Asynchronous event-based high speed vision for microparticle tracking. *Journal of microscopy* **245**(3), 236–244 (2012)
46. Okamoto, K., Nishio, S., Saga, T., Kobayashi, T.: Standard images for particle-image velocimetry. *Measurement Science and Technology* **11**(6), 685 (2000)
47. Pereira, F., Gharib, M., Dabiri, D., Modarress, D.: Defocusing digital particle image velocimetry: a 3-component 3-dimensional DPIV measurement technique. application to bubbly flows. *Experiments in Fluids* **29**(1), S078–S084 (2000)
48. Pereira, F., Gharib, M.: Defocusing digital particle image velocimetry and the three-dimensional characterization of two-phase flows. *Measurement Science and Technology* **13**(5), 683 (2002)
49. Posch, C., Matolin, D., Wohlgenannt, R.: A QVGA 143 dB dynamic range frame-free PWM image sensor with lossless pixel-level video compression and time-domain CDS. *IEEE Journal of Solid-State Circuits* **46**(1), 259–275 (2010)
50. Prasad, A.K.: Stereoscopic particle image velocimetry. *Experiments in Fluids* **29**(2), 103–116 (2000)
51. Raffel, M.: Background-oriented schlieren (BOS) techniques. *Experiments in Fluids* **56**(3), 60 (2015)
52. Raffel, M., Willert, C.E., Scarano, F., Kähler, C.J., Wereley, S.T., Kompenhans, J.: *Particle image velocimetry: a practical guide*. Springer (2018)
53. Rebecq, H., Gallego, G., Mueggler, E., Scaramuzza, D.: EMVS: Event-based multi-view stereo3D reconstruction with an event camera in real-time. *International Journal of Computer Vision* **126**(12), 1394–1414 (2018)
54. Rebecq, H., Gehrig, D., Scaramuzza, D.: ESIM: an open event camera simulator. In: *Conference on Robot Learning*. pp. 969–982 (2018)
55. Rebecq, H., Ranftl, R., Koltun, V., Scaramuzza, D.: High speed and high dynamic range video with an event camera. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2019)
56. Richard, H., Raffel, M.: Principle and applications of the background oriented schlieren (BOS) method. *Measurement Science and Technology* **12**(9), 1576 (2001)
57. Rogister, P., Benosman, R., Ieng, S.H., Lichtsteiner, P., Delbruck, T.: Asynchronous event-based binocular stereo matching. *IEEE Transactions on Neural Networks and Learning Systems* **23**(2), 347–353 (2011)
58. Ruhnau, P., Guetter, C., Putze, T., Schnörr, C.: A variational approach for particle tracking velocimetry. *Measurement Science and Technology* **16**(7), 1449 (2005)
59. Ruhnau, P., Schnörr, C.: Optical stokes flow estimation: an imaging-based control approach. *Experiments in Fluids* **42**(1), 61–78 (2007)

60. Ruhnau, P., Stahl, A., Schnörr, C.: On-line variational estimation of dynamical fluid flows with physics-based spatio-temporal regularization. In: Joint Pattern Recognition Symposium. pp. 444–454. Springer (2006)
61. Scarano, F.: Tomographic PIV: principles and practice. *Measurement Science and Technology* **24**(1), 012001 (2012)
62. Schanz, D., Gesemann, S., Schröder, A.: Shake-The-Box: Lagrangian particle tracking at high particle image densities. *Experiments in fluids* **57**(5), 70 (2016)
63. Schneiders, J.F., Scarano, F.: Dense velocity reconstruction from tomographic PTV with material derivatives. *Experiments in fluids* **57**(9), 139 (2016)
64. Serrano-Gotarredona, T., Linares-Barranco, B.: A 128×128 1.5% contrast sensitivity 0.9% FPN $3 \mu\text{s}$ latency 4 mW asynchronous frame-free dynamic vision sensor using transimpedance preamplifiers. *IEEE Journal of Solid-State Circuits* **48**(3), 827–838 (2013)
65. Stam, J.: Stable fluids. In: Proceedings of the 26th annual conference on Computer graphics and interactive techniques. pp. 121–128 (1999)
66. Tan, Z.P., Thurow, B.S.: Time-resolved 3D flow-measurement with a single plenoptic-camera. In: AIAA Scitech 2019 Forum. p. 0267 (2019)
67. Tropea, C., Yarin, A.L.: Springer handbook of experimental fluid mechanics. Springer Science & Business Media (2007)
68. Wang, H., Liao, M., Zhang, Q., Yang, R., Turk, G.: Physically guided liquid surface modeling from videos. *ACM Transactions on Graphics (TOG)* **28**(3), 1–11 (2009)
69. Xiong, J., Fu, Q., Idoughi, R., Heidrich, W.: Reconfigurable rainbow PIV for 3D flow measurement. In: Proc. ICCP. pp. 1–9. IEEE (2018)
70. 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 Trans. Graph.* **36**(4), 36 (2017)
71. Yoon, S.Y., Kim, K.C.: 3D particle position and 3D velocity field measurement in a microvolume via the defocusing concept. *Measurement Science and Technology* **17**(11), 2897 (2006)
72. Zang, G., Idoughi, R., Wang, C., Bennett, A., Du, J., Skeen, S., Roberts, W.L., Wonka, P., Heidrich, W.: TomoFluid: Reconstructing dynamic fluid from sparse view videos. In: Proc. CVPR. pp. 1870–1879 (2020)
73. Zhu, A.Z., Atanasov, N., Daniilidis, K.: Event-based feature tracking with probabilistic data association. In: 2017 IEEE International Conference on Robotics and Automation (ICRA). pp. 4465–4470. IEEE (2017)
74. Zhu, A.Z., Yuan, L., Chaney, K., Daniilidis, K.: EV-FlowNet: Self-supervised optical flow estimation for event-based cameras. arXiv preprint arXiv:1802.06898 (2018)
75. Zihao Zhu, A., Atanasov, N., Daniilidis, K.: Event-based visual inertial odometry. In: Proc. CVPR. pp. 5391–5399 (2017)