# Multi-Pulse Dual-Frame Velocimetry

M. Moaven[1], D. W. Carter[1,*]

1: Dept. of Mechanical, Materials & Aerospace Engineering, Illinois Institute of Technology, Chicago USA
*Corresponding author: dcarter10@illinoistech.edu

**Keywords:** Particle Tracking Velocimetry, Lagrangian Particle Tracking, Multi-Pulse Velocimetry

ABSTRACT

We present a novel strategy, termed Multi-Pulse Dual-Frame Velocimetry (MPDFV), to track the advection of flow tracers to obtain volumetric Lagrangian Particle Tracks (LPT) with high time-resolution using low-speed dual-frame cameras. A timing scheme is employed that pulses a volume-filling LED once in the first frame and up to seven times in the second frame. By taking advantage of the rapid pulsing of the LED and frame-straddling, the scheme resolves the problem of directional ambiguity while also effectively collecting kHz-resolved data in time. To overcome the large seeding density of particles in the second frame, the "Peak-CNN" machine learning architecture (Godbersen et al., 2024) is used to identify peak locations and the open-source "OpenLPT" code (Tan et al., 2020) is used to identify the particle tracks. Finally, Lagrangian velocities are assimilated to an Eulerian-grid using physics-informed neural networks. MPDFV enables time-resolved PTV with lower-cost higher-resolution dual-frame cameras: improving affordability and data resolution. The framework is demonstrated on two data sets: a direct numerical simulation of homogeneous turbulence (with a ground-truth reference) and a large-scale turbulence box experiment (with comparison to established time-resolved PTV).

## 1. Introduction

Diagnostics of high-speed flows and flows containing strong gradients requires fine temporal sampling (on the order of 1-100µs) to capture and sufficiently characterize particle trajectories. While illumination systems (lasers/LEDs) adopted for particle image velocimetry/particle tracking velocimetry (PIV/PTV) typically operate up to $\sim$100kHz, conventional cameras are limited by the read-out time required to transfer pixel information into the camera's internal memory, capping continuous acquisition to rates on the order of $\sim$1kHz (when operated at full spatial resolution). The most common techniques to enable smaller inter-frame separation times, $dt$, include: (1) reducing the usable region of the camera sensor, or (2) operating in "frame-straddling" mode whereby information from a first exposure is rapidly stored onto an adjacent memory element, allowing acquisition of a second frame at a $dt$ much smaller than typical read-out times. Both approaches

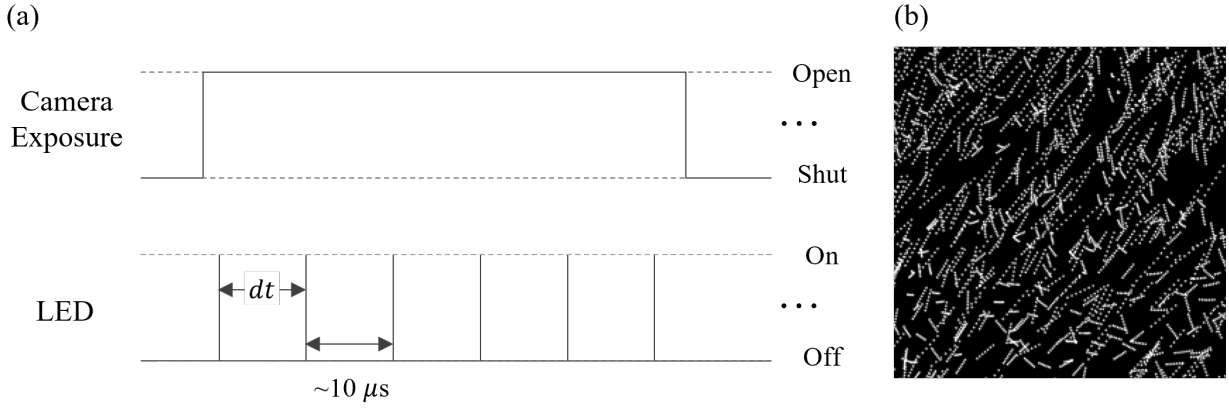



**Figure 1.** Example of a timing diagram (a) and corresponding particle image (b) for multi-pulse velocimetry.

are associated with obvious drawbacks: while the first option lowers the achievable spatial resolution, the second limits temporal information to a single time difference, precluding access to higher-order quantities such as acceleration. These limitations motivate the development of alternative acquisition strategies that can bridge the gap between current imaging and illumination capabilities.

One such approach, particle streak velocimetry (PSV), takes inspiration from primitive forms of PTV, wherein a particle's motion is encoded onto an extended camera exposure as a streak (Zhang et al., 2024). However, ambiguity exists in deciphering temporal content between the endpoints of the streak. A similar approach to PSV involves leveraging the higher frequency capabilities of lasers/LEDs by pulsing the illumination source multiple times within the duration of an exposure. This encodes distinct realizations of the particle field onto a single frame, as illustrated in figure 1. The resulting image provides access to temporal information sampled at much finer rates than conventional time-resolved PIV/PTV. The primary distinction between multi-pulse approaches and PIV/PTV, however, is that the latter explicitly relates each particle realization to its associated time stamp through the frame number, while the former technique involves resolving this ambiguity as a post-processing step. In the case of advection dominated flows, this can be relatively trivial to solve, e.g., through the use of an initial predictor field. However, complications arise when analyzing zero-mean flows with strong gradients and overlapping projections. This is commonly referred to as the problem of directional ambiguity.

A variety of approaches can be taken to address directional ambiguity. An early technique focused on shaping the illumination pulse train to produce a distinct pattern that can be matched to prescribed pulse widths (e.g., dot-streak-dot-dot) (Agüí & Jimenez, 1987). Temporal information has also been encoded in the image intensity map through "tagging" primary pulses with a shorter secondary pulse to generate particle images that are distinct in size and intensity (Grant & Liu, 1990). The relative orientation of the particle images can subsequently be matched to the

order of the prescribed pulse widths to provide directional information. The multi-pulse shake-the-box (MP-STB) algorithm requires solving for direction across a pair of double-pulse images (Novara et al., 2019). This is achieved through bidirectional extrapolation from each image, using a velocity prediction provided by Particle Space Correlation (Novara et al., 2016), to find agreement between intermediate positions. Color-sequence particle streak velocimetry (CSPSV) encodes temporal information through RGB illumination (Wang et al., 2019), enabling direction determination via correspondence to the imposed color sequence, albeit at the cost of requiring color-sensitive imaging hardware. More recently, (Scarano et al., 2025) introduced a method that utilizes asymmetric and irregularly spaced pulse sequences, allowing directionality to be inferred through correlating spatial particle distances with temporal pulse separations. The authors identified seeding density limitations due to the superimposition of multiple exposures onto one image, and the associated formation of ghost particles, as primary areas for improvement. While the aforementioned methods have demonstrated varying degrees of success, we propose an alternative strategy for resolving directional ambiguity that is compatible with existing PIV/PTV hardware and may allow access to higher effective seeding densities.

The proposed approach leverages the small $dt$ available in double-frame imaging to freeze an initial realization of the particle field onto the first frame, before firing multiple illumination pulses during the second exposure, allowing access to finely sampled temporal information. Setting the inter-frame time separation such that the corresponding particle distance is much less than the average particle spacing enables a relatively straightforward track initialization that can be extrapolated forward in time to provide predictions for the remainder of the track. We refer to this method as multi-pulse dual-frame velocimetry (MPDFV). A key challenge associated with multi-pulse approaches is the increased imaging density resulting from exposing a single image to multiple realizations of the same particle field. To address this, we leverage modern techniques in particle identification and flow field reconstruction. The following sections describe the working principle of MPDFV before demonstrating the technique using synthetic numerical experiments as well as experimental measurements.

## 2. Methodology

MPDFV, illustrated in figure 2, combines concepts from frame-straddling and multi-pulse velocimetry. The LED is timed to pulse once in the first frame and multiple times in the second frame. This allows the starting position of a given track to be easily localized from a sparse image (Frame 1), while its propagation through time is embedded into a long exposure image (Frame 2). After triangulating all particle projections in the measurement volume, a 3D tracking algorithm can be implemented to build a temporal correspondence for each track. Data assimilation methods can subsequently be leveraged for a full volumetric reconstruction.

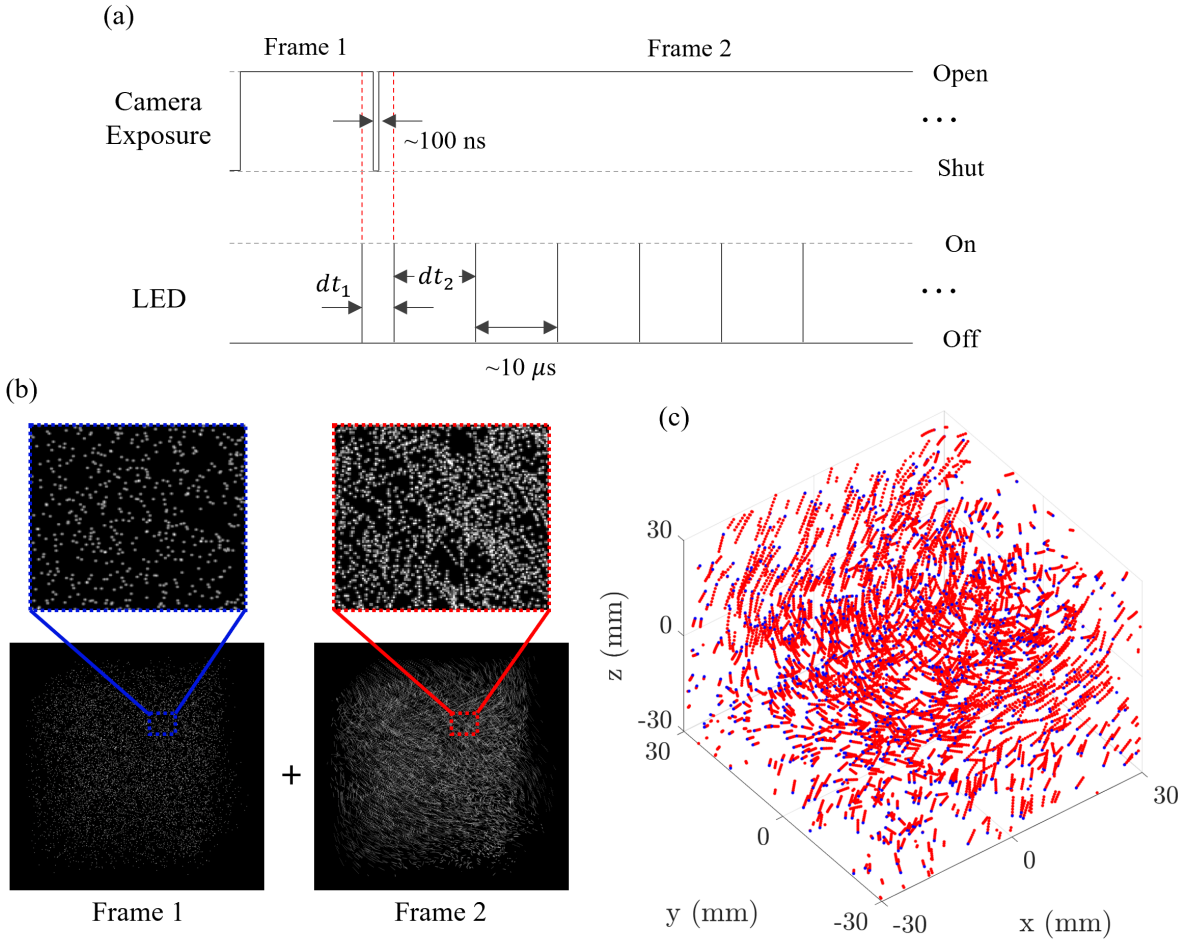



**Figure 2.** Illustration of multi-pulse dual-frame velocimetry. Figure 2(a) demonstrates a typical timing scheme while (b) and (c) show the corresponding image pair and particle field reconstruction, respectively.

The first step of the MPDFV processing pipeline is to identify particle projections across all views. For this, we adopt Peak-CNN – a neural network-based approach that excels in subpixel peak identification in dense images (Godbersen et al., 2024). Rather than relying on intensity thresholds or correlation with predefined Gaussian masks, Peak-CNN leverages a convolutional neural network (CNN) to predict pixels containing a particle image peak before calculating their subpixel center position. The CNN model parameters are optimized using a set of labeled training images that can be generated either synthetically or through stacking together multiple low density experimental images (Godbersen et al., 2024).

Stereomatching across views and subsequent triangulation are performed through partial incorporation of the open-source implementation of shake-the-box (OpenLPT) (Tan et al., 2020). OpenLPT implements a novel two-step approach to removing ghost particles. In the first step, each set of potential matches is assigned a preference vector that counts the number of times a particle's projection is associated with the lowest triangulation error across all sets of candidate matches. The second step prunes through all potential matches and removes candidates that contain duplicate projections associated with a lower preference vector, leading to a globally optimized solution. In a comparison to simple filtering using a triangulation error threshold, the authors showed the percentage of ghost particles dropped from 110% to 26% at 0.125ppp (where "ppp" is the particles per pixel, Tan et al., 2020).

In our current implementation, building a temporal correspondence between particle realizations across several pulses relies on the average particle spacing far exceeding the displacement between Frame 1 and the first pulse belonging to Frame 2. This allows for a simple nearest neighbor search (within a user-defined window) around each particle triangulated from Frame 1 to identify candidates for track initialization. The weighted distance is computed in both 3D as well as each 2D projection, forming a combined cost function:

$$\phi_p = ||\tilde{X}_p^{3D} - X^{3D}||M + ||\tilde{X}_p^{2D} - X^{2D}|| \tag{1}$$

where $M$ refers to the image magnification while $\tilde{X}$ and $X$ denote candidate and reference particles, respectively. The particle, $p$, corresponding to the minimum cost is linked to the reference particle in Frame 1 in order to begin the track. Note that a global prediction can be provided for advection-dominated flows, allowing a reduction in the search region dimensions.

Following initialization, each track is extrapolated forward in time, according to the ratio $dt_2/dt_1$, to estimate the expected position corresponding to the subsequent pulse. The cost function here is modified to consider the distance between candidates to the prediction. Again, the particle associated with the lowest weighted distance is chosen to continue the track. This process is repeated through the entire pulse train and tracks are propagated using predictions provided according to the assumption of minimum change in velocity. If no candidate particles are found in the 2D or 3D search regions, the prediction is stored as a placeholder and continuation of the track is attempted

based on extrapolation to the subsequent pulse. In the event that no viable candidates are found in two consecutive timesteps, the track is terminated. Once particles are linked to a track, they are removed from the pool of available candidates to avoid duplicates. An inherent limitation of MPDFV is that tracking is predicated on the successful identification and triangulation of a particle in Frame 1. Thus, a key consideration for the user is the potential drop in reconstruction quality around the borders of the field of view.

While the tracking procedure builds a spatio-temporal map of discrete particle realizations, MPDFV recovers a 3D reconstruction of the measured velocity field via a physics-informed neural network (PINN). In the context of particle tracking, PINNs have recently emerged as a means to regularize noisy measurements and generate dense velocity fields by incorporating known physical constraints into the reconstruction process (Raissi et al., 2019). This is achieved through minimization of a composite loss function that includes deviations from governing partial differential equations (PDEs). We adopt the PINN implementation presented by Clark Di Leoni et al. (2023), which consists of a loss function that takes the form:

$$L = \lambda_d L_d + \lambda_p L_p \tag{2}$$

where $\lambda$ denotes the weight assigned to each term. $L_d$ and $L_p$ correspond to the data and physics losses, respectively:

$$L_d = \frac{1}{N_u} \sum_{i}^{N_u} \| u_{pinn}(x^i_{data}, t^i_{data}) - u^i_{input} \|_2^2 \tag{3}$$

$$L_p = \frac{1}{N_p} \sum_{i}^{N_p} \sum_{j}^{N_f} \| PDE_j(x^i_p, t^i_p) \|_2^2 \tag{4}$$

$N_u$, $N_p$, and $N_f$ refer to the number of measured velocity points, collocation points, and PDEs, respectively. In the present case, the PDEs are the Navier-Stokes equations for mass and momentum conservation. Note that the minimization of $L_p$ promotes adherence to the governing physics while $L_d$ ensures the output velocity field respects the observed measurements. Adopting a neural network here allows for a continuous representation of the flow field that can be interrogated at any spatio-temporal coordinate. Since the network produces a continuous and differentiable representation of the flow field within the measurement domain, gradient quantities such as the velocity gradient tensor can be recovered without introducing errors due to finite differences. Thus, the output of MPDFV consists of Eulerian fields of velocity/pressure in order to facilitate data visualization/comparison and derivative calculations. In the context of turbulent flows, which have statistics that are notoriously slow to converge, a separate PINN must be obtained for each image pair. This can be a limiting factor in the implementation of MPDFV, as the PINNs are relatively expensive to compute.

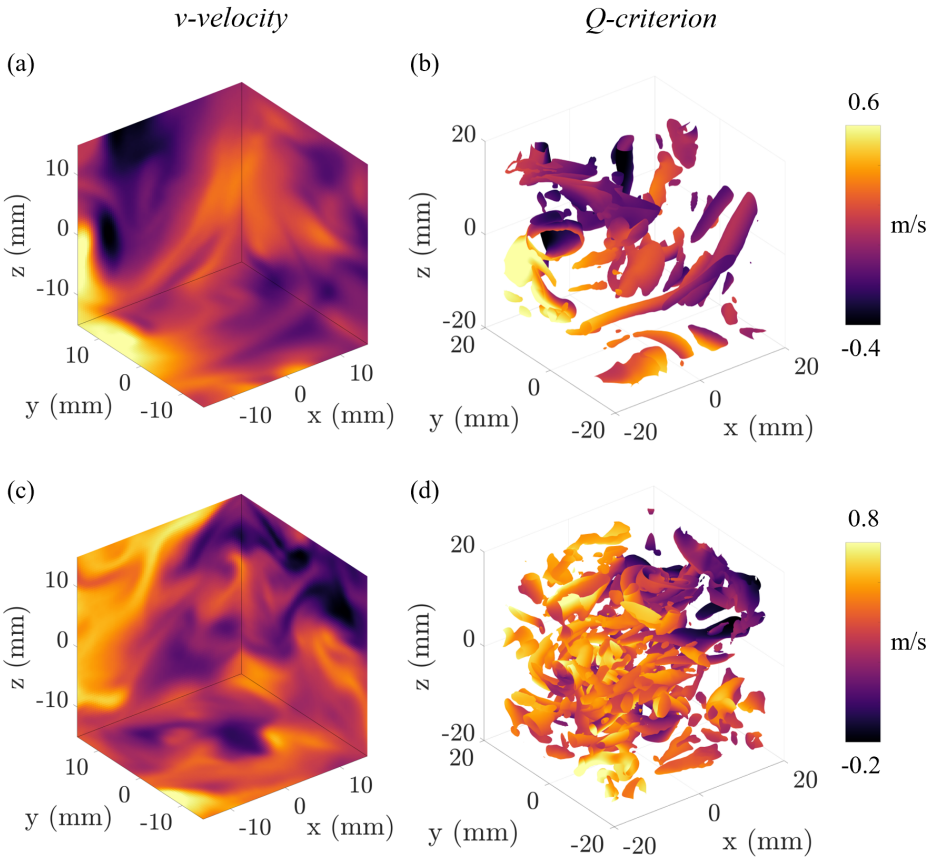



**Figure 3.** Ground truth velocity fields (single-component $v$ velocity shown) of homogeneous isotropic turbulence (left) and corresponding positive Q-criterion iso-surfaces (right) for Vol. 1 (top) and Vol. 2 (bottom) used to validate MPDFV.

## 3. Dataset Description

Similar to other multi-pulse techniques, the performance of MPDFV is expected to be limited by the overlapping of multiple realizations of the particle field – while a higher number of pulses increases the information provided by each track, the resulting increase in imaging density presents more ambiguity in particle identification and triangulation. Thus, the balance between spatial density and pulse frequency is a crucial consideration that is explored using synthetic measurements. In addition, the separation time in the second frame, $dt_2$, should be configured to mitigate the occurrence of subpixel displacements while providing favorable conditions for tracking.

### 3.1. Synthetic Analysis: Direct Numerical Simulation

The synthetic case used to evaluate MPDFV is a direct numerical simulation (DNS) of homogeneous isotropic turbulence (HIT) at a Taylor-scale Reynolds number of 433, available via the Johns

Hopkins Turbulence Database (JHTBD) (Li et al., 2008). A volume of size $0.39 \times 0.39 \times 0.39$ was interrogated from the center of the periodic domain of dimensions $2\pi \times 2\pi \times 2\pi$. Dimensionalization was performed similarly to that described in (Zhou et al., 2024), resulting in a $50 \times 50 \times 50$ mm$^3$ volume, referred to hereforth as Vol. 1 and illustrated in figures 3(a) and 3(b) by velocity contours and positive Q-criterion iso-surfaces, respectively. To verify the generalizability of the technique, a second independent volume (Vol. 2) of the same dimensions centered around $\pi/4$ was generated from the HIT database (figures 3(c) and 3(d)). Synthetic images of the particle fields were produced using an in-house algorithm. Particles of diameter 0.3 mm are projected onto four simulated cameras spaced in a pyramid configuration at $\pm 20^\circ$. The camera specifications include an image resolution of 1024×1024 pixels and 0.01 mm pixel pitch, resulting in a field of view of 102.4×102.4 mm$^2$ at a magnification of 0.1.

The parameter space for the variables most relevant to MPDFV are particle density, measured in particles per pixel (ppp), time between acquisitions (dt), and the number of pulses per frame ($\zeta$) and are presented in Table 1. Note the subscripts F1 and F2 refer to frame 1 and frame 2, respectively. The combination of particle density with pulse quantity leads to an effective imaging density for frame 2. For example, if ppp = 0.01 and $\zeta_{F2}$=7, the actual density of Frame 2 is $0.07ppp$. A total of 27 datasets were processed for Vol. 1 while a smaller selection was explored for Vol. 2.

| **Parameter** | **Value** |
|---|---|
| $ppp$ | 0.005, 0.01, 0.02 |
| $dt_1$ **(**ms**)** | 0.2 |
| $dt_2$ **(**ms**)** | 0.2, 0.4, 0.8 |
| $\zeta_{F1}$ | 1 |
| $\zeta_{F2}$ | 3, 5, 7 |

**Table 1.** Parameter space explored for synthetic evaluation of MPDFV.

### 3.2. Comparative Analysis: Experimental Particle Tracking in a Large-Scale Turbulence Box

Experimental data was collected in a newly constructed turbulence box facility at the Illinois Institute of Technology. The box is a large-scale air-powered quasi-synthetic random jet array similar to the turbulence box reported by Carter et al. (2016). The facility has four inward-facing panels separated by 1.74 m, each with an array of 8 x 12 MAC solenoid valves mounting 2 mm jet orifices fed by 414 kPa of pressurized air for a total of 384 jets. A schematic of the facility is shown in figure 4. Ninety-six jets on each panel are controlled by a Measurement Computing (MCC) USB-DIO96 controller that uses TTL signals to trigger solid-state relays to provide the needed power to turn the jets on and off with sub-millisecond response time. The power is provided by a DIN-mounted 24V power supply on each panel.

The new turbulence box is custom-designed to accommodate state-of-the-art 3D measurements with optical access from four sides. For the velocimetry, a custom system with four high-speed cameras (4 megapixel Phantom V-641) mounting 135 mm lenses and synchronized with a volumetric LED illumination and 0.3 $\mathrm{mm}$ helium-filled soap bubbles (HFSB) as flow tracers. The HFSB seeding required eight separate nozzles to continuously inject the tracers into the facility from the ceiling during the experiment. Based on the depth of field and illumination provided in the region of interest in center of the chamber, a volume of approximately 12 $\mathrm{cm}$ in $x$, 15 $\mathrm{cm}$ in $y$, and 10 $\mathrm{cm}$ in $z$ was obtained, where $y$ is the vertical direction and $x$ and $z$ are the two spanwise directions.

The facility is capable of producing turbulence with specific properties depending on the settings of the algorithm used to generate the flow as well as the magnitude of the air supply pressure. To generate a controlled flow suitable to test the MPDFV methodology a very low level of turbulence was desired such that the time-resolved images from the high-speed cameras highly resolved the particle displacements in time. This allows for testing temporal sub-sampling of the data. To achieve this, a very low pressurization of the box was used and the flow was driven using the standard sunbathing algorithm (Variano & Cowen, 2008; Carter et al., 2016) with a mean jet on time $\mu_{on} = 250$ ms and source fraction $\phi = \mu_{on}/\mu_{off} = 0.125$, resulting in a Reynolds number $Re_\lambda = \dfrac{u'\lambda}{\nu} = 55$, where $u'$ is the root-mean-square fluctuating velocity, $\lambda$ is the Taylor microscale, and $\nu$ is the kinematic viscosity. To a good approximation, the flow was found to be isotropic in the volume of interest.

Two methods are used to provide comparison for particle tracking in the experimental facility. First, we used the established time-resolved method using the open-source OpenLPT (Lagrangian Particle Tracking) code (Tan et al., 2020) with physics-informed neural-implicit data refinement (Zhou et al., 2024) to recover the fine-scale velocity gradients on an Eulerian grid that provide similar comparisons as the synthetic analysis in the DNS. Twenty sequential snapshots were obtained, with nominally 10,000 particle peaks identified in the images (resulting in $ppp \approx 0.0025$). By taking the same data set with one LED pulse per frame and taking the sum of $N$ consecutive images to create an effective frame 2, the MPDFV methodology can be tested and is compared to time-resolved PTV as a function of $N$ as well as the temporal resolution.

## 4. Results

We first consider a synthetic analysis based on the DNS data where the ground-truth information is directly accessible. We subsequently consider an experimental campaign where traditional time-resolved particle tracking was performed at low seeding density and high temporal resolution, allowing for the assessment of the MPDFV methodology.

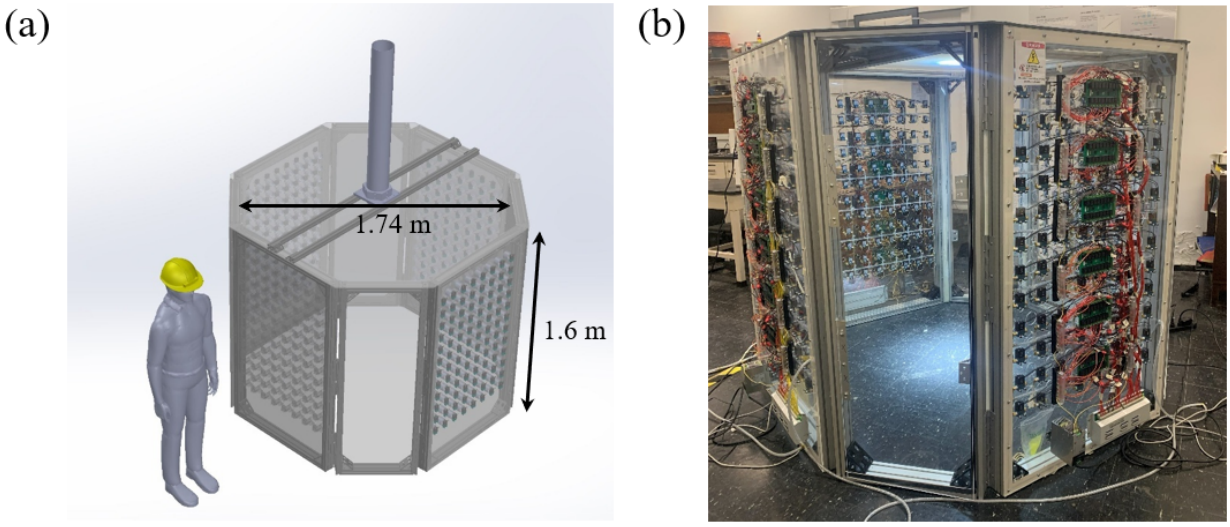



**Figure 4.** CAD rendering of the turbulence box facility (left) and the completed construction in the laboratory (right).

### 4.1. Results of Synthetic DNS Analysis

#### 4.1.1. Localization and Tracking

Figure 5 illustrates the mean localization error and the percentage of correctly tracked particles across the full parameter space for Vol. 1. In all cases, localization errors increase beyond the initial pulse due to the rise in imaging density between frame 1 and frame 2. Similarly, increasing imaging density — either through higher particle seeding or a greater number of illumination pulses — leads to higher average localization errors and a reduction in the associated tracking yield. For example, in Frame 3 at 0.01$ppp$ with $dt_2 = 0.4\text{ms}$ and $\zeta_2 = 3$, the localization error increases by 28.0% when the particle density is doubled to 0.02$ppp$, and by 28.9% and 48.8% in cases wherein the number of pulses in the second frame is increased to $\zeta_2 = 5$ and $\zeta_2 = 7$, respectively. Thus, increases in temporal density have a more pronounced effect on the triangulation error than increases in spatial density, despite resulting in comparable effective imaging densities. This is perhaps to be expected as a rise in pulse frequency guarantees an increase in the local image density for a given particle. In addition, localization errors appear inversely proportional to $dt$. A revealing feature of figure 5 that may explain this relation is the distinct drop in error between the second-to-final pulse and final pulse across all cases representing $dt_2 = 0.2\text{ms}$. As the endpoints of each particle's track are guaranteed to only be neighbored by their own projection on one side, this trend strongly suggests that localization accuracy at lower inter-frame temporal distances is limited by the increased likelihood of consecutive realizations of the same particle overlapping. Conversely, the prediction error grows as the extrapolation distance (i.e., ratio between $dt_2$ and $dt_1$) increases, leading to a decrease in the corresponding tracking yield. This trade-off results in a lower quantity of tracks but with an increase in the average accuracy. The following subsection examines this

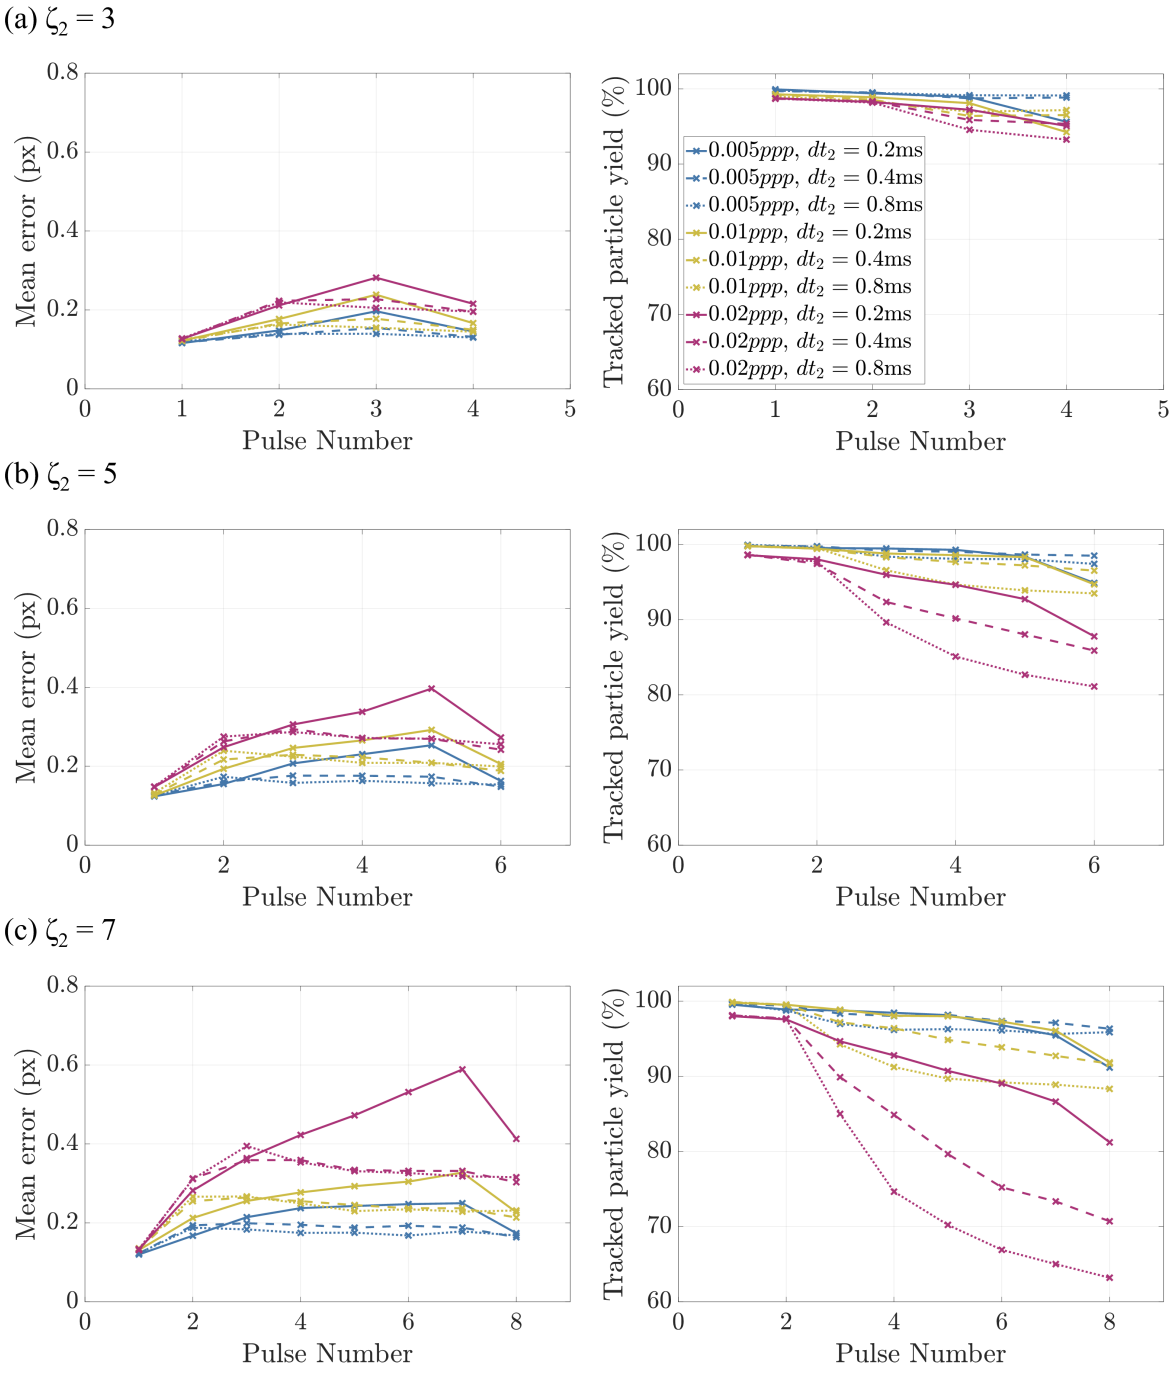



**Figure 5.** Localization (left) and tracking (right) statistics for HIT data processed by MPDFV at each pulse.

balance in terms of its impact on overall reconstruction accuracy.

### 4.1.2. Reconstruction Analysis

The performance of MPDFV in reconstructing 3D vector fields in a turbulent flow field is assessed using the L2 norm of the normalized error, $\epsilon$:

$$\epsilon = \sqrt{\epsilon_u^2 + \epsilon_v^2 + \epsilon_w^2} \tag{5}$$

where

$$\epsilon_u = \frac{|u_{MPDFV} - u_{truth}|}{u'_{truth}} \tag{6}$$

| Imaging density (*ppp*) | $dt_2$ (ms) | $\zeta_2$ | Effective density (*ppp*) | $\mu_\epsilon$ | $\sigma_\epsilon$ |
|---|---|---|---|---|---|
| 0.005 | 0.2 | 3 | 0.015 | 0.140 | 0.101 |
| 0.005 | 0.2 | 5 | 0.025 | 0.172 | 0.131 |
| 0.005 | 0.2 | 7 | 0.035 | 0.199 | 0.143 |
| 0.005 | 0.4 | 3 | 0.015 | 0.117 | 0.081 |
| 0.005 | 0.4 | 5 | 0.025 | 0.117 | 0.095 |
| 0.005 | 0.4 | 7 | 0.035 | 0.132 | 0.106 |
| 0.005 | 0.8 | 3 | 0.015 | 0.098 | 0.080 |
| 0.005 | 0.8 | 5 | 0.025 | 0.092 | 0.076 |
| 0.005 | 0.8 | 7 | 0.035 | 0.087 | 0.070 |
| 0.01 | 0.2 | 3 | 0.03 | 0.132 | 0.090 |
| 0.01 | 0.2 | 5 | 0.05 | 0.163 | 0.122 |
| 0.01 | 0.2 | 7 | 0.07 | 0.203 | 0.151 |
| 0.01 | 0.4 | 3 | 0.03 | 0.089 | 0.077 |
| 0.01 | 0.4 | 5 | 0.05 | 0.111 | 0.089 |
| 0.01 | 0.4 | 7 | 0.07 | 0.123 | 0.105 |
| 0.01 | 0.8 | 3 | 0.03 | 0.072 | 0.063 |
| 0.01 | 0.8 | 5 | 0.05 | 0.074 | 0.060 |
| 0.01 | 0.8 | 7 | 0.07 | 0.074 | 0.066 |
| 0.02 | 0.2 | 3 | 0.06 | 0.123 | 0.092 |
| 0.02 | 0.2 | 5 | 0.1 | 0.177 | 0.125 |
| 0.02 | 0.2 | 7 | 0.14 | 0.246 | 0.165 |
| 0.02 | 0.4 | 3 | 0.06 | 0.087 | 0.068 |
| 0.02 | 0.4 | 5 | 0.1 | 0.099 | 0.085 |
| 0.02 | 0.4 | 7 | 0.14 | 0.148 | 0.127 |
| 0.02 | 0.8 | 3 | 0.06 | 0.065 | 0.055 |
| 0.02 | 0.8 | 5 | 0.1 | 0.065 | 0.056 |
| 0.02 | 0.8 | 7 | 0.14 | 0.113 | 0.107 |

**Table 2.** Average ($\mu_\epsilon$) and standard deviations ($\sigma_\epsilon$) of MPDFV reconstruction errors across parameter sweep for Vol. 1.

Table 2 presents the average, $\mu$, and standard deviation, $\sigma$, of $\epsilon$ across the parameter space of Vol. 1. There exists a strong coupling between $\mu_\epsilon$ and $\sigma_\epsilon$, indicating that the reconstruction quality of the flow field degrades throughout the domain when conditions are made more challenging.

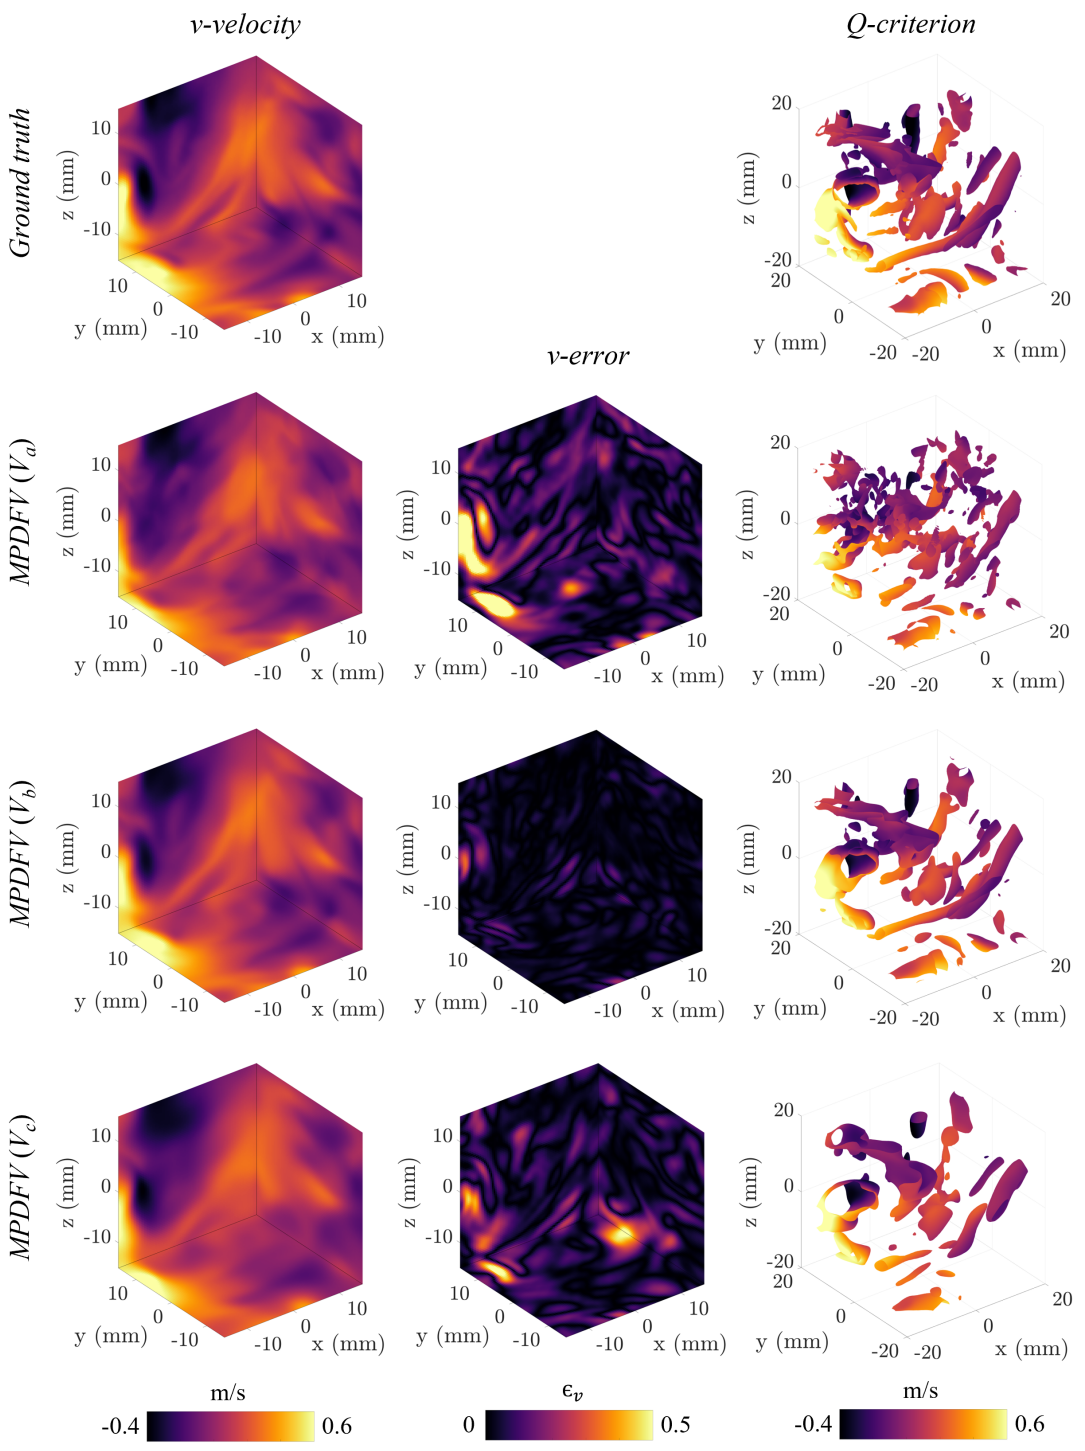



**Figure 6.** Contours of $v$-velocity (left), corresponding error fields (middle), and iso-surfaces of Q-criterion (right) for Vol. 1.

Consistent with the reduction in localization error observed in figure 5, a distinct trend with respect to $dt_2$ exists, wherein larger temporal spacings (and therefore increased distances between particle projections along a track) yield lower errors in the reconstructed velocity field. As the recovery of velocity gradients benefits from a more dense temporal sampling, these results indicate that an improvement in the algorithm's handling of overlapping particle projections should be a primary focus of future work. While figure 5 demonstrated that increasing particle density generally leads to higher localization errors and reduced tracking yield, table 2 reveals a nonlinear dependence with respect to reconstruction error. On average, $\mu_\epsilon$ decreases from 0.128 to 0.116 as the imaging density increases from 0.005 $ppp$ to 0.01 $ppp$, before rising to 0.125 at 0.02 $ppp$. The corresponding standard deviations exhibit a similar trend. Notably, the lowest reconstruction errors are obtained at the highest imaging density of 0.02 $ppp$ ($\zeta_2 = 3$ and $\zeta_2 = 5$). This suggests that, despite the general degradation in triangulation accuracy, the PINN benefits from increased spatial sampling while effectively regularizing noise introduced in earlier processing steps.

In general, reconstruction accuracy deteriorates with increasing $\zeta_2$, except at the lowest imaging density and largest temporal spacing. Here, the error mildly decreases by 11.2% as $\zeta_2$ increases from 3 to 7. In addition, the sensitivity of the error to $\zeta_2$ becomes less pronounced as $dt_2$ is increased. This behavior indicates that when particle overlap is mitigated, either through reduced global sampling or increased temporal spacing, the reconstruction can benefit from additional temporal samples. In contrast however, the reconstruction error improves by 26.5% when doubling the spatial density from 0.005$ppp$ to 0.01$ppp$, highlighting that increasing the spatial sampling is a more effective means of improving accuracy than increasing the number of pulses. In addition, the three highest reconstruction errors occur at $dt_2 = 0.2\text{ms}$ and $\zeta_2 = 7$, indicating that local increases in particle overlap limit the performance of MPDFV considerably more than increases in the global density.

Figure 6 presents velocity field contours for the parameter selections corresponding to the lowest and higher average reconstruction error, as well as the lowest overall density image. These are plotted together with the associated error fields and iso-surfaces of Q-criterion. For convenience we refer to the reconstructions as $V_a$, $V_b$, and $V_c$, respectively. The corresponding imaging settings are:

- $V_a$: 0.02$ppp$, $dt_2 = 0.2\text{ms}$, $\zeta_2 = 7$
- $V_b$: 0.02$ppp$, $dt_2 = 0.8\text{ms}$, $\zeta_2 = 3$
- $V_c$: 0.005$ppp$, $dt_2 = 0.2\text{ms}$, $\zeta_2 = 3$

All three velocity fields show good qualitative agreement to the ground truth, with $V_b$ providing the closest match. $V_a$ exhibits higher error magnitudes than $V_b$ almost universally across the flow field, with the largest discrepancies coinciding with regions of strong gradients in the velocity

field. The Q-criterion fields also show distinct differences. While dominant flow structures are captured in both cases, $V_a$ introduces considerable distortions and spurious features, indicative of the increased noisiness in the particle tracks. This elevated localization error, exacerbated by the reduced tracking yield, likely limits the ability of the PINN to distinguish between noise and true velocity gradients when reconstructing such sparse measurements. $V_c$ also shows a general drop in reconstruction quality relative to $V_b$, particularly near high gradient regions. However, in this case the degradation manifests in a loss of finer features rather than the introduction of noise/distortions. This is expected as the rise in error here is due to measurement sparsity rather than reduced localization accuracy. Overall, MPDFV captures the primary features of this complex turbulent flow with good fidelity, although the results provided here indicate further development focusing on extending the operating range with respect to imaging density could prove very effective towards improving the application of MPDFV on flows with fine-scale features.

| Imaging density (*ppp*) | $dt_2$ (ms) | $\zeta_2$ | Effective density (*ppp*) | $\mu_\epsilon$ | $\sigma_\epsilon$ |
|---|---|---|---|---|---|
| 0.005 | 0.8 | 7 | 0.035 | 0.218 | 0.164 |
| 0.01 | 0.8 | 7 | 0.07 | 0.157 | 0.117 |
| 0.02 | 0.8 | 3 | 0.06 | 0.121 | 0.092 |
| 0.02 | 0.8 | 5 | 0.1 | 0.135 | 0.103 |
| 0.02 | 0.8 | 7 | 0.14 | 0.171 | 0.125 |

**Table 3.** Average ($\mu_\epsilon$) and standard deviations ($\sigma_\epsilon$) of MPDFV reconstruction errors across select cases for Vol. 2.

Table 3 summarizes a subset of cases processed for Vol. 2. Relative to Vol. 1, reconstructions of Vol. 2 exhibit uniformly higher $\mu_\epsilon$ and $\sigma_\epsilon$, reflecting the increased difficulty associated with resolving higher frequency velocity content. In addition, there exists a stronger sensitivity to imaging density. At a fixed $dt_2$ of 0.8ms and $\zeta_2 = 7$, $\mu_\epsilon$ drops by 32.5% as the imaging density doubles from 0.005 $ppp$ to 0.01 $ppp$. For comparison, the corresponding changes in Vol. 1 is just 16.1%, highlighting the heightened sensitivity of the current implementation to imaging conditions in flows with larger Reynolds numbers/finer small-scale features. Improving robustness to such variations will therefore be a priority for future developments.

The recovery of structures in the Q-criterion field shown in figure 7 (corresponding to 0.02 $ppp$, $dt_2 = 0.8$ms, $\zeta_2 = 3$) matches the ground truth well, however, capture of fine-scale structures is limited by the relatively sparse spatial sampling (average particle spacing of $7.5\eta$, where $\eta$ denotes the Kolmogorov length scale). Nevertheless, despite the increase in reconstruction error and flow field complexity, the velocity fields shown in figure 7 demonstrate that dominant flow features are captured well. Thus, the application of MPDFV to Vol. 2 provides encouraging results with respect to the generalization of the algorithm.

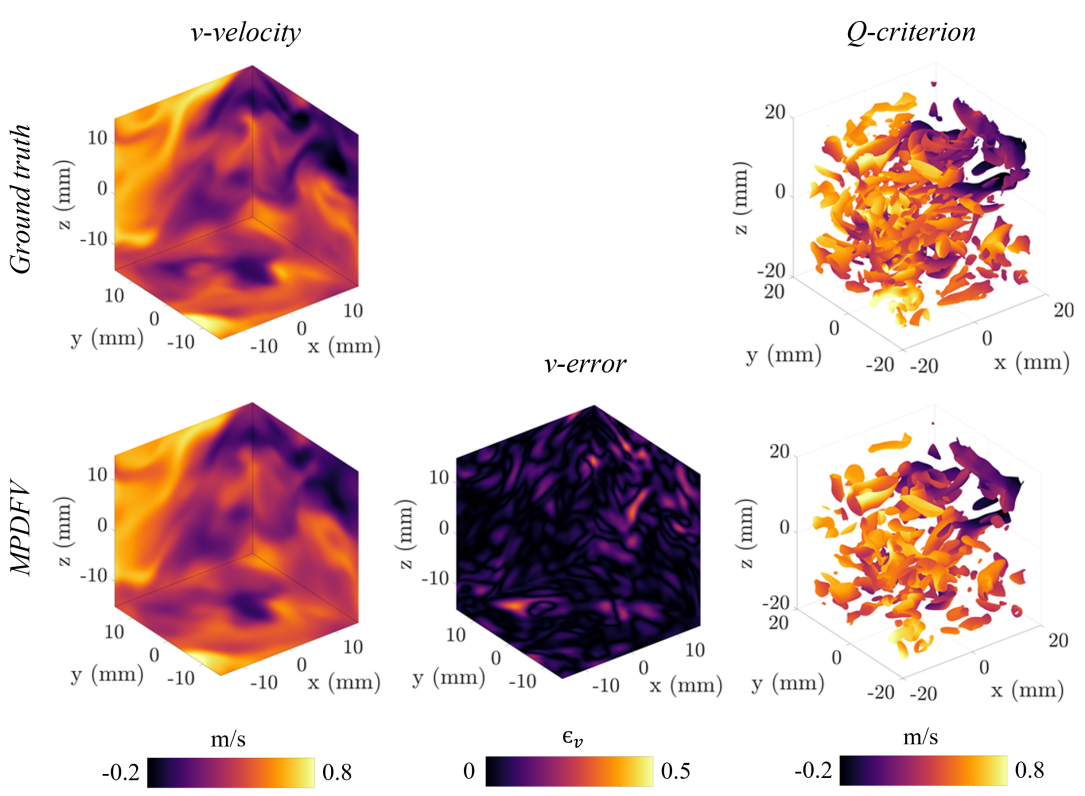



**Figure 7.** Contours of $v$-velocity (left), corresponding error fields (middle), and iso-surfaces of Q-criterion (right) for Vol. 2.

## 4.2. Results of Experimental Campaign

### 4.2.1. Modified Tracking with Shake-the-Box Framework

Three representative runs were selected from the experimental dataset, corresponding to the bottom 5th percentile, median, and top 5th percentile of spatially-averaged strain rate, $\bar{S}_{ij}$, to span the range of flow conditions present in the data, denoted by $\bar{S}_{ij,\,low}$, $\bar{S}_{ij,\,med}$, and $\bar{S}_{ij,\,high}$, respectively. As no ground truth was available for the experimental measurements, PTV-derived tracks were used as the benchmark for comparison. The MPDFV implementation was modified to incorporate particle position refinement and ghost particle suppression via the Iterative Particle Reconstruction/Shake-the-Box (IPR/STB) framework Novara et al. (2016). In this configuration, particle positions are first identified using PeakCNN, after which tracking is performed independently in each projection. The timestamps assigned to each identified peak are then used to construct a pseudo-time-resolved image sequence by isolating particles within a $5\times5$ region centered on the corresponding peak. This allows OpenLPT's STB algorithm to be applied prior to feeding the resulting tracks to the PINN, and provides a more consistent basis for comparison with the PTV measurements, which were also obtained using OpenLPT. Each PTV timestep contained approximately 6,000 particles, while the MPDFV particle count was approximately 70% of this value, as tracking in the latter case is performed while STB remains in its initialization phase.

### 4.2.2. Experimental Tracking Results

Figure 8 compares the reconstructed MPDFV velocity fields with those obtained from PTV. Across all three strain-rate conditions, MPDFV recovers the overall spatial structure of the velocity field and preserves the primary features present in the PTV reconstructions, indicating that the method captures the general flow topology despite the reduced sampling density. Noticeable smoothing of the reconstructed fields is nonetheless evident relative to PTV, particularly in regions of sharp velocity gradients. The absence of any clear correlation between these discrepancies and the mean strain rate suggests that the dominant limitation is associated with sampling density and the reduced effectiveness of track filtering and positional refinement with fewer temporal sampled. The impact of spatial smoothing is reflected more clearly in the probability density functions (PDFs) of one arbitrary diagonal component of the velocity gradient tensor, shown in figure 9. Across all three cases, the MPDFV PDFs exhibit attenuated tails relative to PTV, indicating systematic under prediction of intense gradient events and loss of fine-scale turbulent content in the Eulerian reconstruction. The consistent nature of this tail suppression across strain-rate conditions further points to limited spatial sampling resolution as the primary constraint on reconstruction fidelity, rather than amplification of spurious gradients.

The sensitivity of reconstruction quality to temporal spacing is examined for the medium strain-

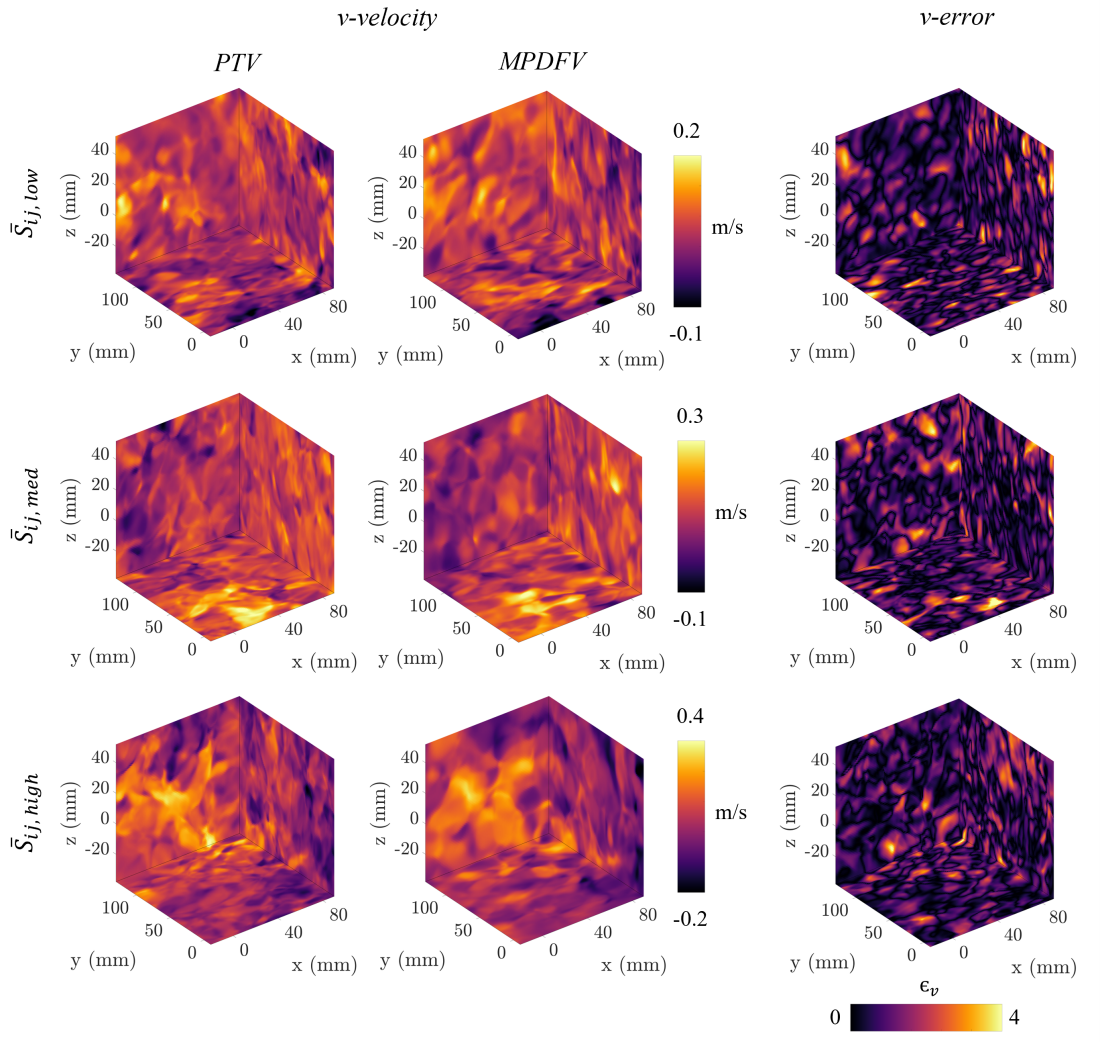



**Figure 8.** Contours of $v$-velocity (left) and corresponding residual fields (right) for experimental homogeneous isotropic turbulence measurements across three averaged strain-rate conditions.

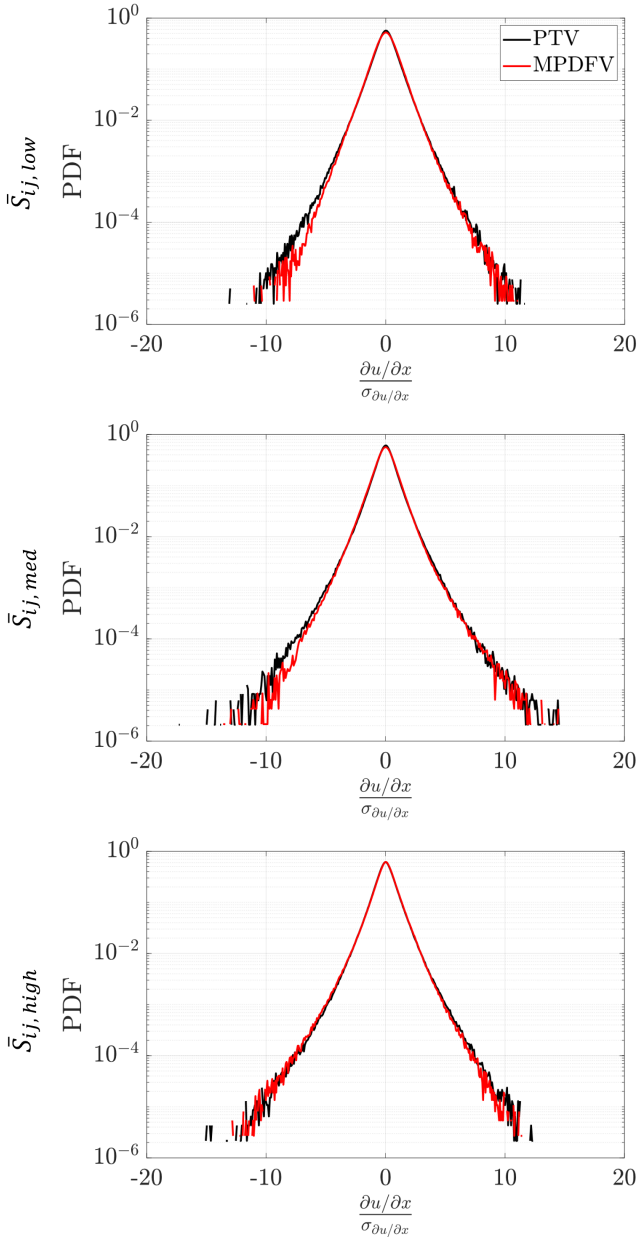



**Figure 9.** Velocity-gradient probability density functions (PDFs) corresponding to the experimental configurations illustrated in figure 8.

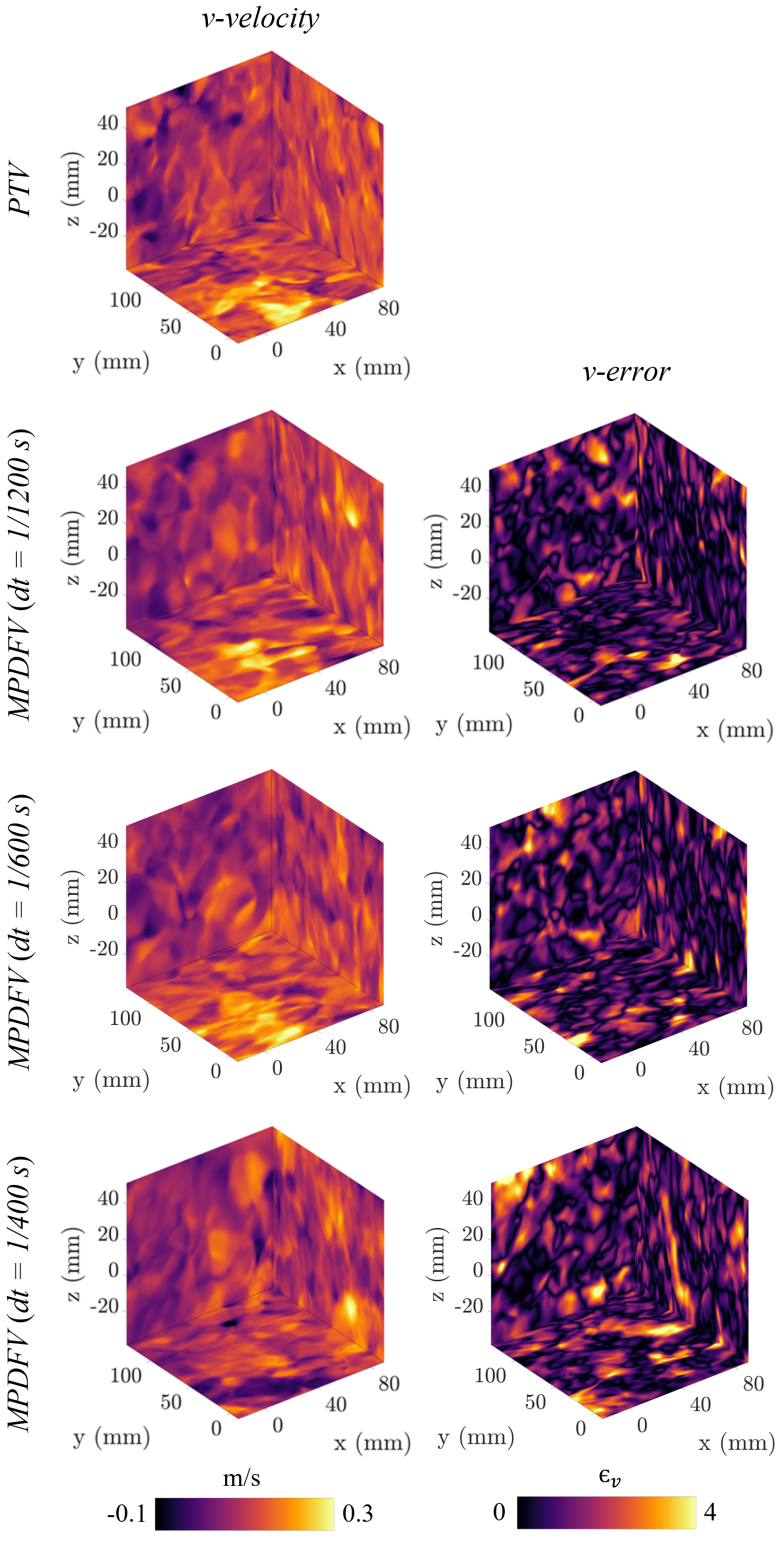



**Figure 10.** Contours of $v$-velocity (left) and corresponding residual fields (right) for experimental homogeneous isotropic turbulence measurements across varying inter-pulse spacing.

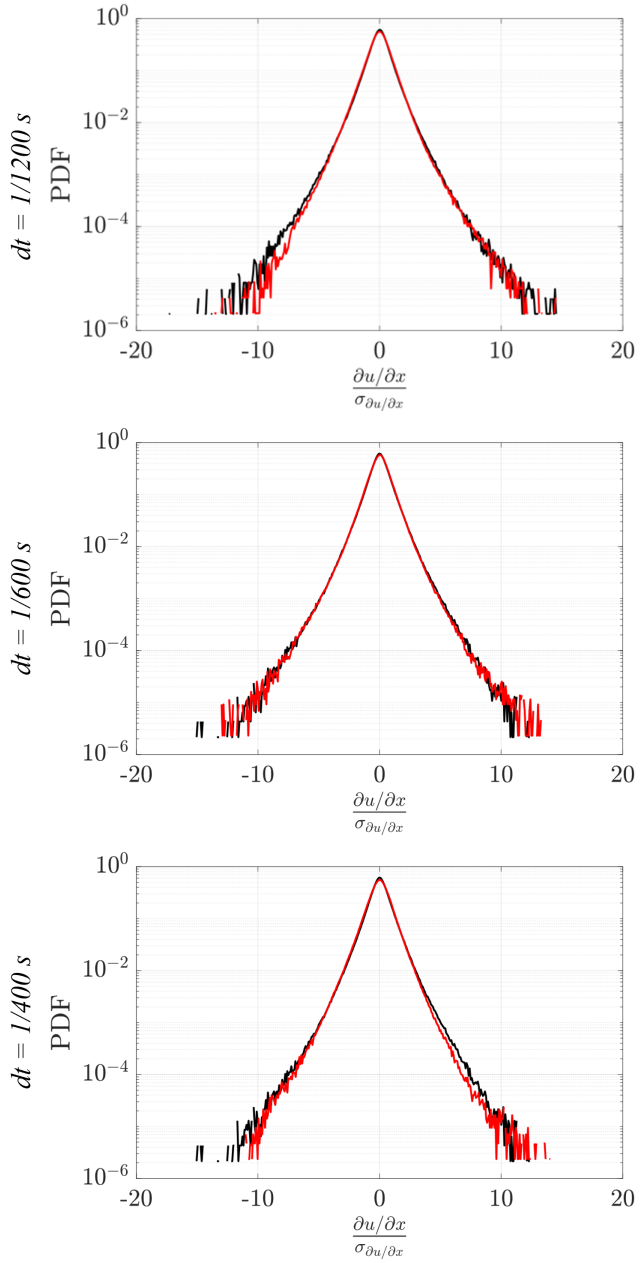



**Figure 11.** Velocity-gradient PDFs corresponding to the experimental configurations illustrated in figure 10.

rate case in figure 10. Note that $dt_1$ and $dt_2$ were kept constant to allow for integration into STB. Unlike the synthetic results, reconstruction quality degrades with increasing $dt$, with the $dt = 1/400s$ case producing a velocity field that bears little resemblance to the PTV measurements. The corresponding error maps show elevated residuals distributed across the entire domain. The PDFs (figure 11), however, remain broadly similar across temporal spacings, suggesting that the degradation stems from track corruption due to ghost particles rather than a systematic bias in the velocity estimation. These findings indicate that temporal spacing is a critical parameter governing reconstruction fidelity, with the optimal value dependent on data quality. For clean synthetic images, the low ghost particle rate means that sub-pixel localization accuracy is improved by ensuring sufficient inter-particle separation. In experimental data, imperfect peak identification generates ghost particles that corrupt tracks and can drive the reconstruction toward an incorrect flow state.

Figure 12 examines the effect of increasing the number of illumination pulses in the second frame ($\zeta_2$ = 3, 5, 7) relative to the PTV benchmark. The large-scale flow features are preserved across all pulse counts; however, increasing $\zeta_2$ produces a progressive loss of spatial detail and a corresponding rise in reconstruction error. Relative to $\zeta_2$ = 3, the $\zeta_2$ = 5 and $\zeta_2$ = 7 reconstructions exhibit increasingly diffuse velocity features and reduced gradient sharpness, consistent with the higher imaging density and localization uncertainty introduced by repeated illumination of the same frame. The degradation is less severe than that observed with increasing dt, likely because the STB implementation benefits from the additional temporal samples. The velocity-gradient PDFs (figure 13) corroborate this trend: as $\zeta_2$ increases, the distributions narrow systematically, with progressive attenuation of the tails relative to PTV. The effect is most pronounced at $\zeta_2$ = 7, where deviations from the benchmark become apparent at relatively modest normalized gradient magnitudes. Taken together, the spatial smoothing in the velocity fields and the narrowing of the gradient distributions indicate that beyond a threshold pulse count, the increase in imaging density becomes the dominant factor limiting reconstruction fidelity in the present implementation, outweighing the benefit of additional temporal information.

Overall, these results demonstrate that while the MPDFV algorithm is capable of reconstructing the general flow topology, further refinements to the implementation must be investigated in order to improve the recovery of fine-scale flow features and an emphasis of future work should be on investigating the relationship between temporal spacing and reconstruction quality.

## 5. Conclusions

A novel velocimetry strategy, termed Multi-Pulse Dual-Frame Velocimetry (MPDFV), has been presented for obtaining time-resolved volumetric Lagrangian particle tracks using low-speed dual-frame cameras. The method combines frame-straddling with multi-pulse illumination, pulsing a volume-filling LED once in the first frame and multiple times in the second, resolving the prob-

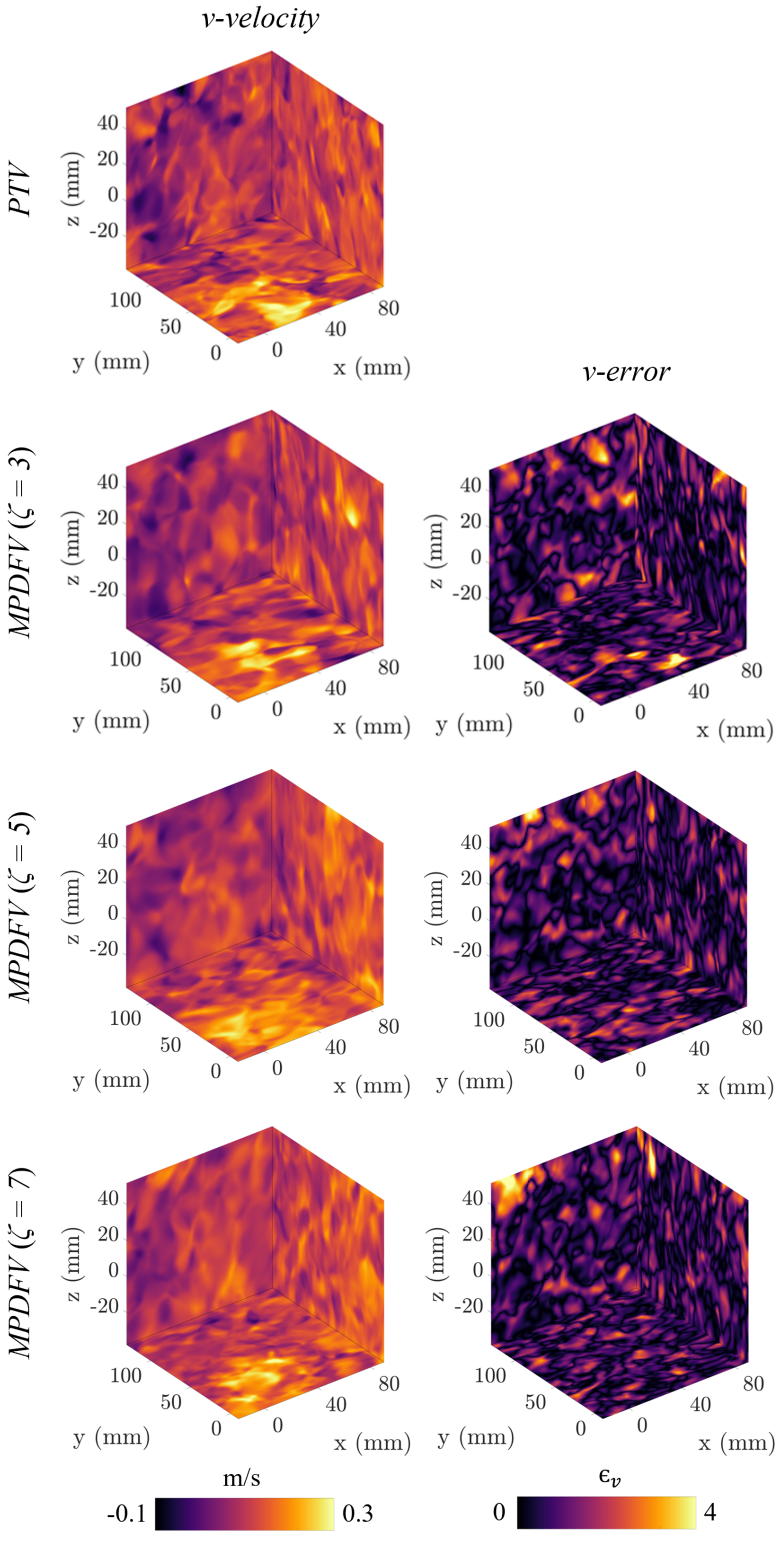



**Figure 12.** Contours of $v$-velocity (left) and corresponding residual fields (right) for experimental homogeneous isotropic turbulence measurements across a varying number of second-frame illumination pulses.

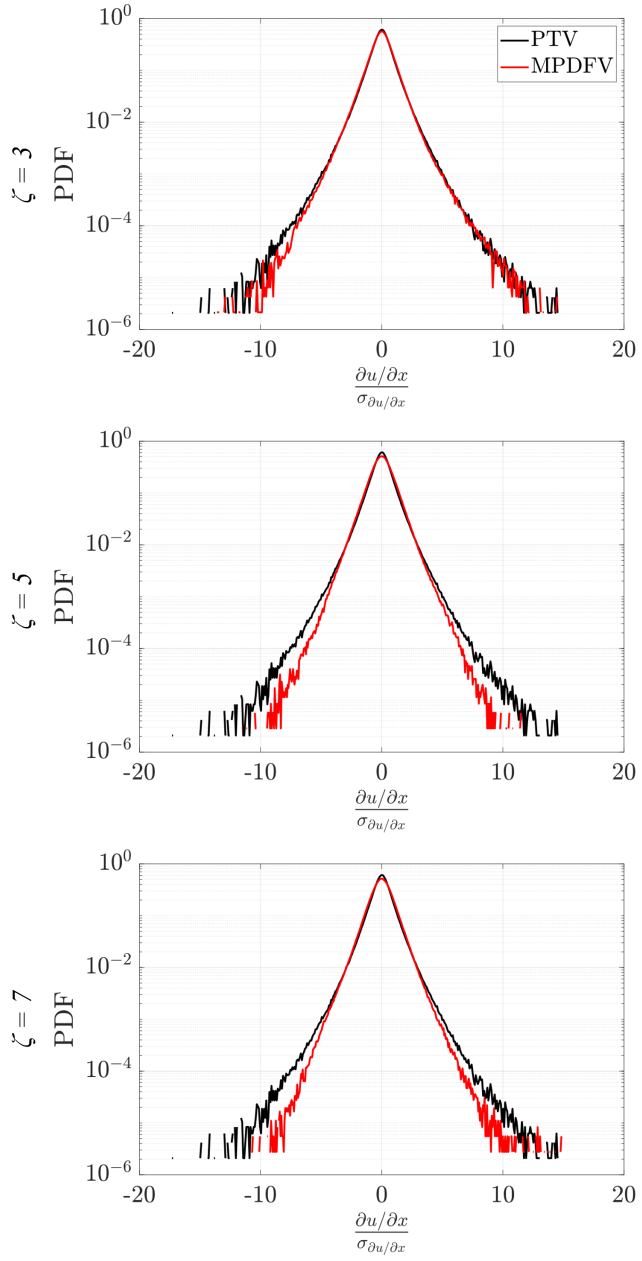



**Figure 13.** Velocity-gradient PDFs corresponding to the experimental configurations illustrated in figure 12.

lem of directional ambiguity while enabling kHz-resolved temporal sampling. The incorporation of PeakCNN allows for particle identification in closely-packed multi-pulse images of a particle's temporal history. Velocity field reconstruction via a physics-informed neural network yields continuous and differentiable Eulerian fields of velocity, pressure, and their gradients. The performance of MPDFV was assessed across a systematic parameter sweep using synthetic images generated from a DNS of homogeneous isotropic turbulence at $Re_\lambda = 433$, and subsequently validated against time-resolved PTV measurements in a large-scale turbulence box experiment.

Synthetic analysis across a wide parameter space verified the feasibility of MPDFV, while highlighting its sensitivity to different imaging conditions. Results suggested a general drop in reconstruction quality with an increase in the number of pulses in the second frame, as the rise in effective imaging density manifests in heightened tracking ambiguity. Localization accuracy at small inter-frame separations was limited by overlapping of consecutive particle realizations, with average errors dropping as $\zeta_2$ increased. Increasing the spatial density from 0.005 $ppp$ to 0.01 $ppp$ yielded a 26.5% reduction in mean reconstruction error, representing a more effective strategy to improving accuracy than increasing the number of illumination pulses. This is due to the rise in local imaging density becoming the dominant limiting factor, outweighing the benefit of additional temporal samples. The lowest reconstruction errors were obtained at the highest spatial density with $\zeta_2 = 3$ and the largest tested temporal spacing.

In the experimental campaign, MPDFV successfully recovered the general topology of the turbulent velocity field across three different strain-rate conditions when compared against PTV. The increase in error with rise in $\zeta_2$ manifested primarily as spatial smoothing in the Eulerian velocity reconstruction and attenuation of the tails of the velocity-gradients PDFs, indicative of limited spatial sampling resolution rather than noise amplification. This suggests that measurement errors caused by a rise in local imaging density tend to result in loss of tracks rather than formation of false tracks. Temporal spacing was identified as a particularly critical parameter in the experimental measurements. Due to the abundance of ghost particles in real measurements, larger $dt$ led to an increased likelihood of imperfect tracking, significantly outweighing the benefit of separating particles in their 2D projections to improve sub-pixel estimation.

These results demonstrate that MPDFV is capable of reconstructing 3D turbulent velocity fields from dual-frame cameras with sufficient fidelity to capture large-scale flow features. The primary limitations of the current implementation are associated with the handling of overlapping particle projections at high imaging densities and short temporal spacings as well as measurement noise and its resulting effect on tracking performance. Future work will focus on improving the robustness of the peak identification and tracking schemes, aiming to extend the operating range of the current algorithm. In addition, a systematic investigation into the effects of experimental noise and its relationship to the optimal temporal spacing is required to improve the generalizability of the technique and provide users with guidance when designing experiments.

**References**

Agüí, J. C., & Jimenez, J. (1987). On the performance of particle tracking. *Journal of fluid mechanics*, *185*, 447–468.

Carter, D., Petersen, A., Amili, O., & Coletti, F. (2016). Generating and controlling homogeneous air turbulence using random jet arrays. *Experiments in Fluids*, *57*(12), 189.

Clark Di Leoni, P., Agarwal, K., Zaki, T. A., Meneveau, C., & Katz, J. (2023). Reconstructing turbulent velocity and pressure fields from under-resolved noisy particle tracks using physics-informed neural networks. *Experiments in Fluids*, *64*(5), 95.

Godbersen, P., Schanz, D., & Schröder, A. (2024). Peak-cnn: improved particle image localization using single-stage cnns. *Experiments in Fluids*, *65*(10), 153.

Grant, I., & Liu, A. (1990). Directional ambiguity resolution in particle image velocimetry by pulse tagging. *Experiments in Fluids*, *10*(2), 71–76.

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

Novara, M., Schanz, D., Geisler, R., Gesemann, S., Voss, C., & Schröder, A. (2019). Multi-exposed recordings for 3d lagrangian particle tracking with multi-pulse shake-the-box. *Experiments in Fluids*, *60*(3), 44.

Novara, M., Schanz, D., Gesemann, S., Lynch, K., & Schröder, A. (2016). Lagrangian 3d particle tracking for multi-pulse systems: performance assessment and application of shake-the-box.

Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, *378*, 686–707.

Scarano, F., Hysa, I., Grille Guerra, A., Tuinstra, M., & Sciacchitano, A. (2025). Asymmetric time sequence for multiple-exposure 3d ptv. *Experiments in Fluids*, *66*(4), 74.

Tan, S., Salibindla, A., Masuk, A. U. M., & Ni, R. (2020). Introducing openlpt: new method of removing ghost particles and high-concentration particle shadow tracking. *Experiments in Fluids*, *61*(2), 47.

Variano, E. A., & Cowen, E. A. (2008). A random-jet-stirred turbulence tank. *Journal of Fluid Mechanics*, *604*, 1–32.

Wang, H., Zhang, H., Hu, X., Luo, M., Wang, G., Li, X., & Zhu, Y. (2019). Measurement of airflow pattern induced by ceiling fan with quad-view colour sequence particle streak velocimetry. *Building and Environment*, *152*, 122–134.

Zhang, D., Tropea, C., Zhou, W., Cai, T., Huang, H., Dong, X., … Cai, X. (2024). Particle streak velocimetry: a review. *Experiments in Fluids*, *65*(9), 130.

Zhou, K., Grauer, S., Schanz, D., Godbersen, P., Schröder, A., Rockstroh, T., … Wieneke, B. (2024). Benchmarking data assimilation algorithms for 3d lagrangian particle tracking. In *21st international symposium on the application of laser and imaging techniques to fluid mechanics, lisbon, portugal* (pp. 1–22).