

# Benchmarking Data Assimilation Algorithms for 3D Lagrangian Particle Tracking

K. Zhou<sup>1</sup>, S.J. Grauer<sup>1,\*</sup>, D. Schanz<sup>2</sup>, P. Godbersen<sup>2</sup>, A. Schröder<sup>2,3</sup>, T. Rockstroh<sup>4</sup>, Y.J. Jeon<sup>4</sup>, B. Wieneke<sup>4</sup>

1: Dept. of Mechanical Engineering, Pennsylvania State University, United States

2: Inst. of Aerodynamics and Flow Technology, German Aerospace Center (DLR), Germany

3: Inst. of Transport Technology, Dept. of Image Based Measurement Techniques, Brandenburgische Technische Universität, Germany

4: LaVision GmbH, Germany

\*Corresponding author: [sgrauer@psu.edu](mailto:sgrauer@psu.edu)

**Keywords:** Lagrangian particle tracking, particle tracking velocimetry, flow reconstruction, data assimilation, turbulence.

## ABSTRACT

This paper reports a comparison of three state-of-the-art data assimilation (DA) algorithms for 3D Lagrangian particle tracking (LPT): Vortex-In-Cell sharp (VIC#), FlowFit3, and neural-implicit particle advection (NIPA). Particle trajectories, termed “tracks,” are spatially sparse, and a reconstruction algorithm is commonly employed to estimate dense Eulerian fields using position, velocity, and/or acceleration data from the tracks. DA algorithms for LPT combine the tracks with a set of governing equations (or constraints) to enhance the accuracy of the velocity field, relative to interpolation methods, and infer additional quantities like pressure. We compare the performance of VIC#, FlowFit3, and NIPA with respect to their accuracy and computational cost. Accuracy is assessed in terms of the mean inter-particle distance, frame rate, and magnitude of tracking errors, and costs are reported in wall time. Synthetic and experimental data sets for homogeneous isotropic turbulence and turbulent boundary layer flows are evaluated; direct numerical simulations are used for the synthetic tests, and realistic localization errors are added to the simulated tracks. All three DA methods perform well in cases with the highest spatio-temporal sampling of the flow, and all three methods are relatively robust to the frame rate and magnitude of localization errors. NIPA is more resilient to sparse seeding than FlowFit or VIC#, which exhibit similar performance, but NIPA is also the most expensive technique by some margin. DA reconstructions are superior to interpolation across the board, with a reduction of error up to 80% in the experimental demonstration.

---

## 1. Introduction

Time-resolved, 3D velocity and pressure fields are widely used to study turbulent flows (Scarano, 2012; Schröder & Schanz, 2023). For instance, 3D velocity data are needed to resolve vorticity and capture coherent structures, and temporal resolution can be employed to infer pressure and determine unsteady aerodynamic loads on a complex body. Particle-based velocimetry is a family of techniques that can deliver spatially-resolved measurements of fluid flows. Of these methods, Lagrangian particle tracking (LPT), also called particle tracking velocimetry, converts time-resolved particle images from several views into 3D trajectories for a large number of tracer particles in the measurement volume (Schröder & Schanz, 2023). Compared to tomographic particle image

velocimetry (tomo-PIV), another 3D velocimetry technique (Scarano, 2012), LPT achieves a higher spatial resolution with minimal spatial filtering and ghost particles. In LPT, individual tracer particles are imaged, localized, and traced, resulting in a set of Lagrangian trajectories that are called “tracks.” Advanced LPT techniques, like Shake-The-Box (STB) (Schanz et al., 2016), perform particle localization and tracking synergistically to deliver tracks that can spatio-temporally resolve turbulent structures. Many applications call for Eulerian velocity, pressure, or other fields to compute derivatives, extract coherent structures, integrate loads, etc. This need gives rise to a flow reconstruction problem, i.e., deducing continuous Eulerian fields from the Lagrangian tracks. The most successful reconstruction algorithms utilize data assimilation (DA) to combine LPT tracks with physics-based equations to mitigate errors in the data and infer latent fields like pressure or density. This paper compares three such algorithms: *Vortex-In-Cell sharp* (VIC#), *FlowFit3*, and *neural-implicit particle advection* (NIPA). Results of our study may help to inform LPT users about best practices and guide the development of next-generation algorithms.

Most early methods for converting scattered tracks into gridded Eulerian data used interpolation. Examples include polynomial (Malik & Dracos, 1995; Imaichi & Ohmi, 1983), natural neighbour (Miozzi, 2005), and Kriging (Agterberg, 2014) interpolation as well as adaptive Gaussian windowing (Agüí & Jimenez, 1987). While these naïve schemes are simple and fast, the resulting estimates are susceptible to noise and prone to over-smoothing. Moreover, the resolution of interpolated fields is limited by the Nyquist frequency of particle sampling in space and time (Stüer & Blaser, 2000). Physics-based information has thus been incorporated into flow reconstruction algorithms to enhance the estimates’ accuracy and resolution, which amounts to an LPT DA problem.

Modern DA algorithms for LPT seek a representation of the flow that simultaneously matches the particle tracks and satisfies physical constraints, usually in a weak sense (Casa & Krueger, 2013; Sperotto et al., 2022; Li & Pan, 2024). Flow states are discretized with an appropriate basis (or similar), like radial basis functions (RBFs) (Casa & Krueger, 2013), B-splines (Gesemann et al., 2016), or empirical functions from a flow decomposition (H. Wang et al., 2015); coefficients of the basis functions are estimated by minimizing an objective loss that comprises “data” and “physics” terms. The former term compares the estimated flow states to measured data, and the latter term indicates deviations from solutions to selected physical constraints, which may include the Navier–Stokes, Euler, or vorticity transport equations. Optimized flow states are roughly consistent with both the tracks and the constraints. Popular interpolation-based DA algorithms for incompressible flows include the FlowFit (Gesemann et al., 2016; Godbersen et al., 2024) and VIC (Schneiders & Scarano, 2016; Jeon et al., 2022) families, which use a B-spline and RBF basis, respectively. Recent progress in machine learning (ML) algorithms and hardware have enabled a spate of ML-based DA algorithms for LPT, with physics-informed neural networks (PINNs) as a major example (Raissi et al., 2019; Zhou et al., 2023; Cai et al., 2024). Instead of using a conventional basis, PINNs represent the flow with one or more neural networks that map space–time input coordinates to flow fields at that location. This function is trained with an objective loss in the same manner as FlowFit, VIC, and other LPT DA algorithms. PINNs are easy to implement and can use the compressible Navier–Stokes and Euler equations for DA with minimal information about the

boundary conditions (Molnar & Grauer, 2022; Molnar et al., 2023).

Given the flurry of developments in LPT DA algorithms, there is a need to benchmark their accuracy and computational cost. Several related efforts have been reported in the recent past. For instance, van Gent et al. (2017) compared the accuracy of pressure reconstructions based on PIV and LPT measurements, including an assessment of VIC+ and FlowFit2, among other algorithms. The authors found that using LPT instead of tomo-PIV leads to more accurate reconstructions due to the availability of Lagrangian velocity and acceleration information from the tracks as well as the lack of spatial filtering. More recently, Sciacchitano et al. (2021, 2022) tested DA methods using particle tracks from a turbulent cylinder wake flow simulation. Reconstructions errors were scored and energy spectra were analyzed under different seeding densities to understand the ability of DA methods to provide reliable information between the tracks. Analysis showed that DA increased the resolvable length scale by a factor of four compared to linear interpolation. However, this work only involved one synthetic flow case and did not consider common sources of experimental error.

This paper is an extension of the aforementioned campaigns that includes up-to-date DA algorithms for LPT. We report a comprehensive assessment of LPT DA methods using multiple numerical and experimental data sets. Mature algorithms like FlowFit3 (Godbersen et al., 2024) and VIC# (Jeon et al., 2022) are compared to a newly-developed ML method, NIPA, which solves equations of motion for the particles as well as the fluid (Zhou et al., 2023; Zhou & Grauer, 2023). The accuracy and cost of each method are assessed with synthetic tests. Several factors are varied in these tests, including the mean inter-particle distance, temporal resolution, and magnitude of tracking errors. Experimental data sets are also used to evaluate the algorithms under practical, real-world conditions. Performance of each DA method is analyzed in terms of errors and computational cost to provide insight into the limits of reconstruction algorithms for turbulent flows.

## 2. DA algorithms for LPT

Three DA algorithms are evaluated in this benchmarking campaign: *FlowFit* from the German Aerospace Center (DLR), *VIC#* from LaVision GmbH, and *NIPA* from the Pennsylvania State University. These algorithms are summarized, below, followed by a description of a binning method that we use to baseline the DA techniques.

### 2.1. FlowFit

FlowFit represents instantaneous flow fields using a 3D B-spline basis, which is fit to particle data obtained from the TrackFit algorithm (Gesemann et al., 2016). B-splines may be efficiently evaluated at arbitrary points in the domain, allowing for easy computation of the fields and their spatial derivatives at selected locations. Coefficients for the basis functions are determined by minimizing the track residuals at the particle locations together with residuals of the incompressible Navier-Stokes equations using a gradient-based optimizer. The coefficients are initialized with a rough first-order approximation. FlowFit is currently a purely 3D method that is independently applied to each time step of a measurement, using velocities and material derivatives from TrackFit. A

key feature of the current version, FlowFit3 (Godbersen et al., 2024), is the improved handling of mass continuity. Instead of including the continuity equation as an additional term in the physics loss, as done by previous instantiations of FlowFit, null-divergence is imposed as a hard constraint by restricting the optimizer to search a divergence-free subspace of velocity fields. This approach prevents the optimizer from exploring many non-physical states and speeds up convergence.

There are two modes of FlowFit, termed “linear” and “nonlinear,” hereafter referred to as FF-linear and FF-nonlinear. The linear mode enforces mass continuity and uses the velocity of particles for optimization. The LPT data loss is combined with a regularization term that penalizes curvature of the velocity field. FF-linear is very fast but does not compute pressure fields. By contrast, the nonlinear mode includes the momentum equation as well as the particle acceleration data from TrackFit. The flow is parameterized using coefficients for the velocity field, only, and the pressure field is computed by a fast-Fourier-transform pressure Poisson solver. This mode not only recovers the pressure field, but it also enhances the accuracy of velocity estimates by including additional physics and acceleration data from the tracks. Consequently, FF-nonlinear comes at a higher computational cost than FF-linear.

Both variants of FlowFit are optimized using the limited-memory Broyden–Fletcher–Goldfarb–Shanno (L-BFGS) algorithm, which is terminated once the objective loss plateaus. Parameters are kept relatively consistent across the test cases considered in this paper. We use a B-spline grid resolution of 0.2 particles-per-cell (ppc), which leads to a coarsening of the grid for the sparser cases, decreasing computational effort. This optimization is disabled for the HIT case, in which we use a fixed grid size for all cases, where the grid corresponds to 0.2 ppc in the densest seeding case. All cases use a curvature penalization factor of 0.03, which is slightly adjusted in a few of the extreme tests based on visual inspection.

## 2.2. VIC#

Vortex-In-Cell algorithms are based on the method proposed by Christiansen (1973). Schneiders & Scarano (2016) created the VIC+ algorithm, extending Christiansen’s method to enable DA for particle-based measurements, and Jeon et al. (2022) implemented a number of improvements to VIC+, culminating in VIC#. The main idea of VIC algorithms is to represent the flow in terms of the vorticity field, which is discretized using RBFs. A sequence of Poisson equations is solved to obtain the velocity and pressure fields, which are the primary quantities of interest. VIC# features several key improvements over VIC+. Additional boundary conditions have been implemented to better obey continuity and the Navier–Stokes equations, and a coarse-to-fine, multi-resolution technique has been implemented to speed up convergence and mitigate dependencies on the grid initialization.

Vorticity RBF coefficients are determined by minimizing a cost function that comprises a series of loss terms. A full description of this function can be found in Jeon et al. (2022). The first two terms are data losses based on the measured velocity and acceleration of the particles, which are

obtained from the tracking algorithm;\* these terms are included in VIC+ as well as VIC#. The latter method, VIC#, includes additional losses. To start, the divergence of vorticity, velocity, and the material derivative of velocity are penalized. Further penalties are applied to the irrotational and solenoidal components of the pressure field. The vorticity field and its material derivative are converted to velocity and its derivative by solving two Poisson equations, and pressure is obtained from velocity via a third Poisson equation. All three of these equations are solved using an algebraic multigrid method (Stüben, 2001). Derivatives of the penalty terms with respect to the RBF coefficients are computed using the adjoint approach of Giering & Kaminski (1998), and the cost function is minimized by an L-BFGS solver.

The basic version of VIC# is a pure 3D solver, denoted VIC#-3D, and only uses temporal derivatives from the particle tracks, akin to FF-linear. Jeon (2023) recently extended VIC# to perform time-resolved reconstructions by introducing a vorticity time-marching method, resulting in the VIC#-4D algorithm. In the 4D variant, the vorticity transport equation is solved using temporal derivatives from the RBF-based estimate. Time marching in VIC#-4D is conducted using a second-order Euler scheme. For each flow considered in this work, the VIC# grid size is determined using the case with the highest seeding density.

### 2.3. NIPA

Neural-implicit particle advection uses PINNs to recover flow states from unfiltered, noisy LPT tracks (Zhou et al., 2023; Zhou & Grauer, 2023; Grauer & Zhou, 2024). NIPA has two modules: neural flow states and data-constraint tracks (DCTs), which represent the flow fields and particle states, in turn. Flow states are parameterized by one or more neural networks, and it is common to use a Fourier encoding to combat the low-frequency bias of gradient-descent-type optimizers (S. Wang et al., 2023). Partial derivatives of the flow with respect to  $\mathbf{x}$  and  $t$ , obtained by automatic differentiation, are used to evaluate the governing equations; the equation residuals are integrated over the flow domain and measurement interval and added up to form a “flow physics loss.” Both the compressible and incompressible Navier–Stokes equations were used to form this loss in Zhou & Grauer (2023). A DCT is a trainable function, e.g., a high-order polynomial or similar, that describes a particle’s velocity over time. The function is constrained to integrate to the tracked particle locations. Integration and differentiation of the DCT yield the particle’s position and acceleration history, respectively. A “particle physics loss” is computed by adding integrated residuals from a suitable equation of motion, such as the Maxey–Riley equation, for each particle. Simultaneously minimizing the flow and particle physics losses yields flow states that approximately solve the governing equations, even when there is particle slip (Zhou & Grauer, 2023).

Some LPT setups, including single-camera systems based on plenoptic imaging or digital in-line holography, suffer from large, anisotropic localization errors. NIPA can be configured to estimate the particle positions alongside the flow states in these cases. Position variables in the DCTs are included as trainable parameters in the optimization, and a “data loss” is added to the flow and particle physics losses. The data loss may be based on weighted residuals in 3D space, or it

---

\*We use the Savitzky–Golay (Savitzky & Golay, 1964) method to filter the tracks via second-order polynomials.

can be formulated in terms of image residuals. Optimizing all three losses, together, amounts to synchronous flow reconstruction and physics-enhanced particle tracking. Grauer & Zhou (2024) provide a detailed description of particle localization via NIPA.

In this work, two deep, feed-forward neural networks are used to represent each flow: one network for velocity and another for pressure. The velocity networks have ten hidden layers with 300 neurons apiece, and the pressure networks have ten hidden layers with 150 neurons. Both networks are augmented with a Fourier encoding to enhance expressivity (Tancik et al., 2020). Unfiltered tracks that span the measurement interval are fed to NIPA. The tracks are split into 11-position segments to form the DCTs, per Zhou & Grauer (2023). Polynomials up to the 14th order are used to represent each component of the particles' velocity. The algorithm is implemented in TensorFlow 2.10, and training is conducted using the Adam optimizer. The initial learning rate of  $10^{-3}$  is lowered to  $10^{-4}$  after the loss plateaus. Training persists to convergence in all cases, which takes around 50,000 epochs at each learning rate.

## 2.4. Binning

Binning is a simple and computationally cheap alternative to the DA schemes described above. In this work, we divide the domain into interrogation volumes, i.e., "bins," and collect all the particles in each bin at a given instance. Binned velocities are a weighted second-order polynomial fit of the scattered velocities inside a bin (Agüera et al., 2016). We use a Gaussian weighting function, which peaks at the bin center (Agüí & Jimenez, 1987). Velocities at the particle positions are obtained by Savitzky–Golay filtering, utilizing second-order polynomials with a filter width of nine time steps, and the velocity field is iteratively filtered to minimize divergence. Pressure is computed from the filtered velocity field by solving a Poisson equation. Grids used for binning are the same as the grids used by the VIC# algorithm.

## 3. LPT measurement scenarios

Two synthetic cases and one experimental case are employed for DA benchmarking. In our numerical tests, we vary the seeding density, temporal resolutions, and magnitude of localization errors to investigate the effect of these factors on flow reconstruction. Experimental data are included to evaluate the performance of LPT DA techniques using real-world measurements.

### 3.1. Synthetic cases

#### 3.1.1. Homogeneous isotropic turbulence

The first synthetic case uses a direct numerical simulation (DNS) of homogeneous isotropic turbulence (HIT) from the Johns Hopkins Turbulence Database (JHTDB) (Perlman et al., 2007). The flow has a Taylor microscale Reynolds number of 433 and is simulated in a periodic domain of size  $2\pi \times 2\pi \times 2\pi$ , discretized into  $1024^3$  voxels. A forcing term is included to compensate for energy dissipation (Rosales & Meneveau, 2005). The data are stored at a time separation of 0.002. Our region of interest (RoI) is the central  $128^3$ -voxel subvolume, and our test spans 201 frames. Assuming air with a kinematic viscosity of  $15 \text{ mm}^2/\text{s}$ , the dimensionalized DNS domain has integral

length, time, and velocity scales of  $l_0 = 127.32$  mm,  $t_0 = 0.2$  s, and  $u_0 = 636.81$  mm/s, respectively. Table 1 lists the key dimensional parameters of both the HIT and boundary layer cases, where  $l_\eta$  and  $\tau_\eta$  denote the Kolmogorov length and time scales.

**Table 1.** Dimensional parameters of the synthetic HIT and TBL flow cases.

Case	$Re$	Domain, mm <sup>3</sup>	Time step, ms	$l_\eta$ or $l_\nu$ , mm	$\tau_\eta$ or $\tau_\nu$ , ms
HIT	433	$100 \times 100 \times 100$	0.4	0.35	8.2
TBL	1000	$56 \times 12 \times 4.25$	0.028	0.036	0.09

### 3.1.2. Turbulent boundary layer

The second case involves DNS of a turbulent boundary layer (TBL) in channel flow, which we also obtain from the JHTDB. Driven by a favorable pressure gradient of  $dp/dx = -0.0025$ , fully-developed boundary layers with a friction-velocity Reynolds number up to 1000 form along the top and bottom walls. The full DNS domain spans  $8\pi \times 3\pi \times 2$  and is discretized using  $3072 \times 2304 \times 512$  voxels along the streamwise, spanwise, and vertical directions. The data are saved at a dimensionless rate of 0.0065. Our synthetic test uses 201 frames of a  $126 \times 54 \times 80$ -voxel subvolume, located along the bottom wall. This volume covers  $1546 \times 331 \times 117$  viscous units, which roughly spans the viscous sub- and buffer layers. Once again assuming the fluid to be air, the dimensional domain has integral length, time, and velocity scales of  $l_0 = 36$  mm,  $t_0 = 0.0044$  s, and  $u_0 = 8.3$  m/s, respectively. Dimensional parameters of the TBL are summarized in Table 1, where  $l_\nu$  and  $\tau_\nu$  denote the viscous length and time scales.

### 3.1.3. Simulating LPT tracks

Ground truth tracks are generated by simulating a dense set of particle trajectories in each case. These tracks may be downselected, downsampled in time, and/or corrupted by noise to vary the measurement scenario. To start, massless particles are uniformly distributed in the domain and advected at the local flow velocity to determine their position over time. Advection is performed using a fourth-order Runge–Kutta scheme, subject to periodic boundary conditions. To ensure statistical convergence and avoid boundary effects, the tracks are simulated in an enlarged domain, which is  $\sim 15\%$  bigger than the RoI in each direction, but only particles within the RoI are retained in the data sets. Tracks with fewer than five positions are discarded, which is common practice in tracking algorithms.

Around 100,000 and 280,000 particles are simulated for the HIT and TBL cases, respectively. These seeding densities correspond to an image density of 0.1 and 0.28 particles-per-pixel (ppp) for a 1 MP camera, the latter exceeding the upper bound of STB’s range of operability (Schanz et al., 2016, 2022). Mean particle spacings are  $6l_\eta$  for the HIT flow and  $6l_\nu$  for the TBL, which are sufficient for near-scale-resolving LPT reconstructions, irrespective of the DA algorithm. The temporal resolution of our tracks is half that of the corresponding DNS data, i.e., 0.8 ms for the HIT flow and 0.057 ms for the TBL. The domain size and temporal resolution of synthetic LPT tests are compa-

able to those of real LPT experiments (Schröder et al., 2015; Schröder & Schanz, 2023). Moreover, the TBL case is designed to closely match the experimental case described in Sec. 3.2.

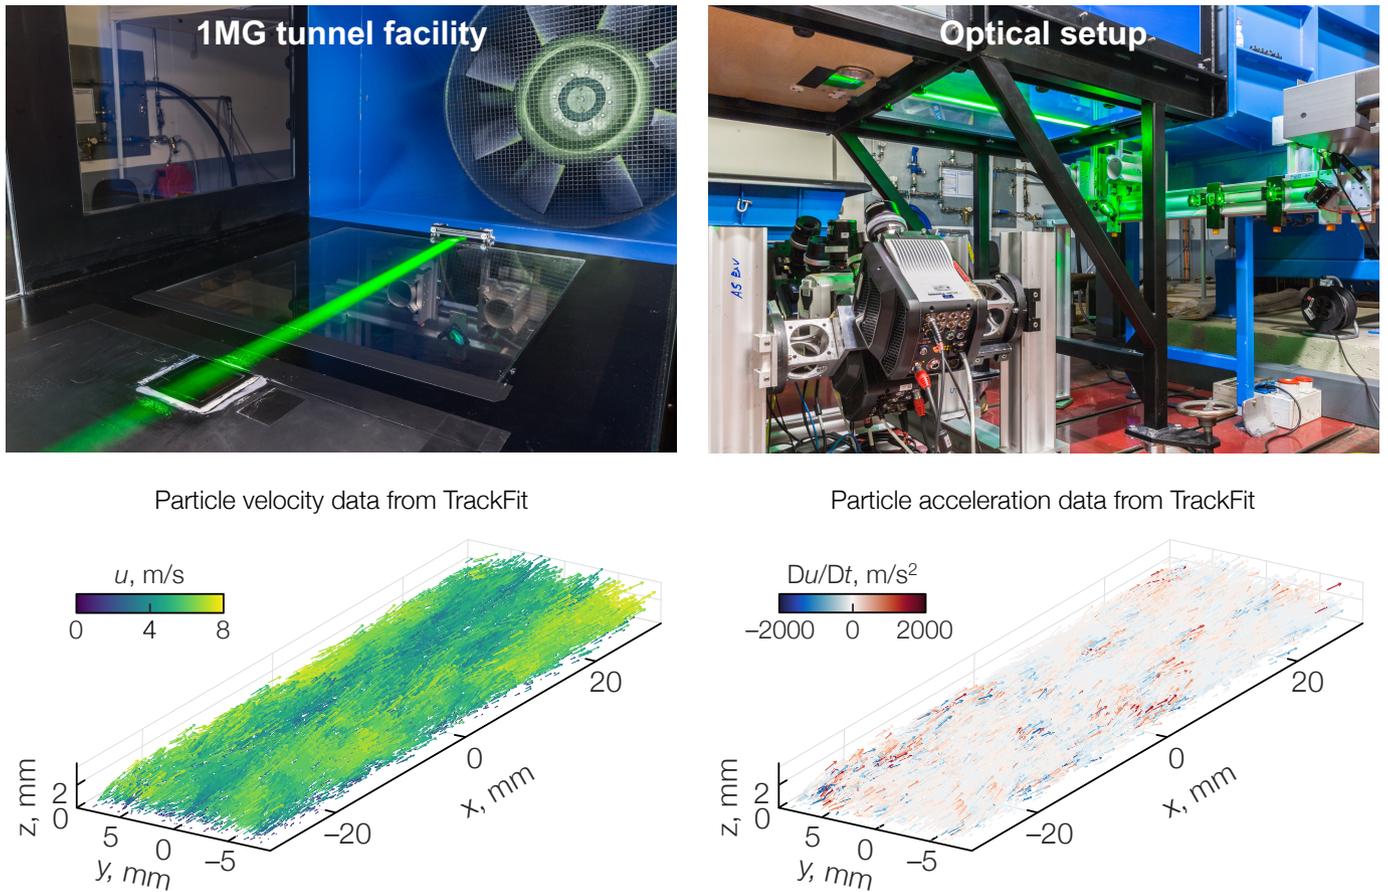
Given the base cases described above, we generate additional data sets by varying the seeding density, frame rate, and localization errors, as detailed below.

1. **Spatial sampling:** we randomly downselect the pool of tracks by a factor of  $k_x = 2^n$ , for  $n = 1, 2, \dots, 6$ , which increases the particle spacing to roughly  $24l_\eta$  and  $24l_\nu$  for the HIT and TBL cases, respectively.
2. **Temporal sampling:** we fix the seeding density to  $k_x = 4$  and then downsample the time steps by a factor of  $k_t = 2, 4, 6, 8$ , thereby decreasing the frame rate of the data set.
3. **Localization errors:** we fix the seeding density and frame rate using  $k_x = 4$  and  $k_t = 1$ , and we corrupt the particle positions with additive, zero-mean Gaussian errors in each direction. The standard deviation of errors in the  $x$ - and  $y$ -directions are equal,  $\sigma_x = \sigma_y = 0.01$  mm and  $0.0033$  mm for the HIT and TBL cases, respectively. These values correspond to a  $0.1$  px uncertainty, assuming a  $1$  MP camera (Schanz et al., 2016; Schröder et al., 2024). Uncertainty in the  $z$ -direction is varied:  $\sigma_z/\sigma_x = 2, 4, 6, 8, 10$ , where, loosely speaking, the optical axis of the imaging system is aligned with the  $z$ -axis. These distributions mimic the anisotropic localization errors that crop up in many LPT measurements, particularly when using a single-camera setup (Sciacchitano et al., 2022; Fahringer et al., 2015; Toloui et al., 2017).

### 3.2. Experimental case

Measurements of a zero-pressure-gradient TBL were performed in DLR's One-Metre Low-Speed Wind Tunnel, Göttingen (1MG). Figure 1 depicts the facility and a sample set of tracks. The 1MG tunnel has a cross-section of  $0.74 \times 1$  m<sup>2</sup> and test section that is  $3$  m long. In our test, airflow entered the test section at a free stream velocity of  $u_\infty = 10$  m/s and encountered a trip, leading to the formation of a TBL along the lower wall. Measurements were made  $2.08$  m downstream of the trip location. The TBL had a friction-velocity Reynolds number of  $995$ , with viscous scales of  $l_\nu = 0.037$  mm and  $\tau_\nu = 0.0925$  ms.

Dioctyl sebacate particles with a mean diameter of approximately  $1.5$   $\mu\text{m}$  were generated by a Laskin-nozzle seeder and homogeneously distributed inside the tunnel. Particles were illuminated within a  $60 \times 15 \times 2.5$  mm<sup>3</sup> RoI (streamwise by spanwise by wall-normal direction) by a dual-head nanosecond laser (Photonics Industries, DM200). Light scattered by the particles was imaged by five high-speed cameras (four Phantom v2640, one v1840) at a repetition rate of  $23.9$  kHz and scaling factor of  $34.185$  px/mm. The camera calibration was initialized using a 3D calibration plate and refined by volume self-calibration (Wieneke, 2008); an optical transfer function was used for iterative particle extraction (Schanz et al., 2012). Images of the flow had a particle image density of around  $0.06$ – $0.07$  ppp. STB was used to track  $50,000$  to  $60,000$  particles over long sequences of up to  $140,000$  consecutive images. Further details of the experimental setup and results can be found in Schröder et al. (2024).



**Figure 1.** Experimental setup at the 1MG tunnel at DLR Göttingen. Top left: Laser-illuminated RoI. Top right: System of five high-speed cameras viewing the RoI from below. Bottom left: instantaneous 3D particle cloud, colored by streamwise velocity. Bottom right: instantaneous 3D particle cloud, colored by streamwise acceleration.

For the present benchmarking campaign, a set of 51 consecutive time steps was randomly chosen, containing 85,285 tracks, in total, with an average of around 54,000 particles per frame. Reconstructions are compared at the 26th time step. To judge the quality of reconstructions, we randomly divide the particles into two groups, one for training and another for testing. The *training set* contains 80% of the particles, which are used to reconstruct the flow; the remaining 20% of particles (in the *test set*) are unseen by the DA algorithms and used to gauge the reconstruction accuracy (Schneiders & Sciacchitano, 2017). The *test set* is further downsampled by a factor of  $k_x = 2^n$ , for  $n = 1, 2, 3$  in subsequent evaluations, akin to the synthetic test cases. This means the DA algorithms utilize between 80% and 10% of the tracked particles for reconstruction, corresponding to mean inter-particle spacings of  $9.4l_v$  and  $18.7l_v$ . Reconstructions are examined within a  $52 \times 12 \times 2.5 \text{ mm}^3$  region that is centered inside the measurement domain, i.e., slightly trimming each border in the wall-normal direction. Velocity fields from the DA algorithms are sampled at the test-set particle positions and compared to the particles' measured velocity.

## 4. Flow reconstruction results

### 4.1. Error metric

Errors are quantified in terms of the normalized root mean square error (NRMSE) of the fluctuating component of each field. This metric corresponds to the energy of errors normalized by the turbulent energy of the field. To illustrate, let  $\varphi$  stand in for the velocity vector field,  $\mathbf{u}$ , or scalar pressure field,  $p$ . We apply the standard Reynolds decomposition,  $\varphi = \bar{\varphi} + \varphi'$ , and compute NRMSEs as follows:

$$e_\varphi = \left( \frac{\langle \|\varphi' - \varphi'_{\text{exact}}\|_2^2 \rangle}{\langle \|\varphi'_{\text{exact}}\|_2^2 \rangle} \right)^{1/2}, \quad (1)$$

where  $e_\varphi$  is the NRMSE of  $\varphi$ ,  $\varphi_{\text{exact}}$  is a ground truth quantity,  $\|\cdot\|_2$  is the Euclidean norm, and  $\langle \cdot \rangle$  is a volumetric average over the RoI,  $\Omega$ ,

$$\langle \varphi \rangle = \frac{1}{|\Omega|} \int_{\Omega} \varphi(\mathbf{x}) \, d\mathbf{x}. \quad (2)$$

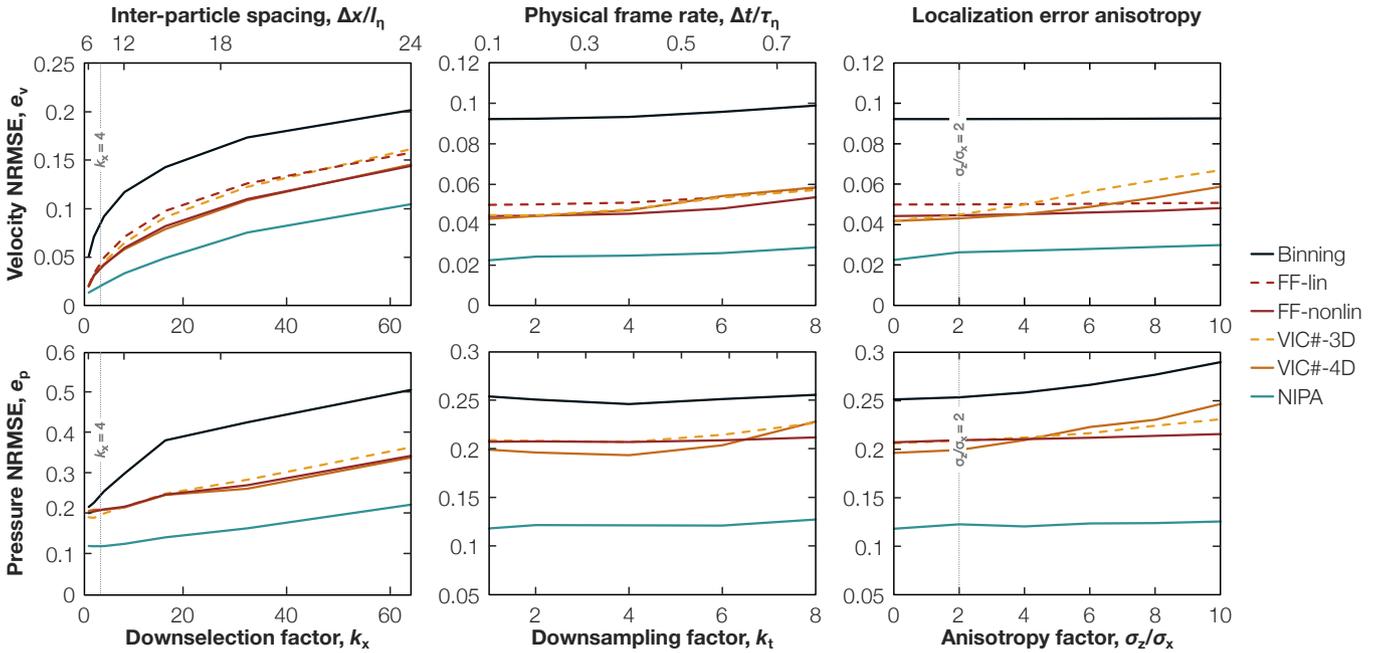
Equation (2) is approximated by summing over elements of the corresponding VIC# grid.

### 4.2. Synthetic HIT results

We first present results from the HIT flow case. Figure 2 shows NRMSEs for velocity and pressure at the central frame; results are shown for binning and all three DA methods under varying mean inter-particle distances, frame rates, and error ratios. The range of spatial and temporal resolution and noise considered here are representative of real-world LPT tests.

As expected, all the reconstructions deteriorate as the particle density decreases, frame rate decreases, and noise increases. The DA methods are consistently more accurate than binning, which indicates the value of incorporating physics into the flow reconstruction algorithm. All the reconstruction techniques are sensitive to spatial resolution (mean inter-particle distance) but robust to temporal resolution and noise. Variants of FlowFit and VIC# exhibit similar errors, although errors in VIC# reconstructions are magnified at large time separations and noise ratios. This effect is probably due to a sub-optimal selection of track filtering parameters. By contrast, NIPA exhibits lower errors in velocity and pressure across the board. This may be attributable to a more effective optimization of the physics loss and the low-frequency inductive bias (smoothing) inherent to the training algorithm (S. Wang et al., 2022). It should be noted that NIPA included the *artificial* forcing term used in the HIT DNS (Rosales & Meneveau, 2005), which does not apply to real flows and was neglected in FlowFit and VIC#. To analyze this difference, we performed another NIPA reconstruction without forcing, using the  $\sigma_z/\sigma_x = 2$  data set. While velocity errors were unchanged, the pressure NRMSE increased from 12.3% to 20.3%, comparable to the FlowFit (20.9%) and VIC# (19.9%) results. This result speaks to the importance of using appropriate governing equations in the inversion (though it does not support the use of synthetic forcing, of course).

Figure 3 depicts cut plots of the HIT velocity and pressure fields at the 26th frame ( $k_x = 4$  and  $\sigma_z/\sigma_x = 2$ ). Results for FF-nonlinear, VIC#-4D, and NIPA are compared to the DNS ground truth.



**Figure 2.** Synthetic HIT velocity and pressure NRMSEs at different conditions for the DA and binning methods. NIPA includes the artificial forcing term, which accounts for the significantly lower pressure error. Vertical dashed lines indicate the conditions corresponding to Fig. 3.

Other variants are omitted due to limited space. Cuts are shown at three orthogonal planes: a left ( $y = 0.4$  mm), rear ( $x = 99.6$  mm), and bottom ( $z = 0.4$  mm) plane, all close to the boundaries of the RoI. DA methods typically exhibit larger errors towards the boundaries due to the unknown in- and outflow conditions. All three DA methods are able to recover turbulent velocity and pressure fluctuations that closely resemble the DNS fields. Quantitatively,  $e_v$  is 4.4%, 4.3%, and 2.6% for FF-nonlinear, VIC#-4D, and NIPA, respectively, and  $e_p$  is 20.9%, 19.9%, and 12.3% (20.3% without the forcing term) in the same order. Absolute error fields are shown on the right side of Fig. 3. The magnitude of apparent errors across these fields is consistent with the trends in Fig. 2.

We further examine the accuracy of reconstructed flow derivatives for the purpose of vortex identification. Figure 4 shows isosurfaces of the ground truth and reconstructed  $Q$ -criterion fields. We compute  $Q$ -criterion values using exact derivatives of the continuous FlowFit and NIPA estimates; second-order finite-differences are used to determine the field for the DNS, binning, and VIC# results. Reconstructions are arranged by method, column-wise, and by seeding density, row-wise. All three DA methods outperform binning at each point in the test matrix. At the highest seeding density, the DA methods resolve the finest coherent structures, all resembling the ground truth well. Vortical structures quickly fade out with sparser seeding: almost all the real structures vanish at  $k_x = 8$ . The DA estimates retain some bona fide structure information, compared to the binning results. Among the DA methods, FF-nonlinear and VIC#-4D exhibit slightly more noise with increasing  $k_x$ , while NIPA produces smoother coherent structures. This result matches the error trends in Figs. 2 and 3. However, whether such smoothing is physical or not requires further examination, e.g., in terms of velocity error spectra.

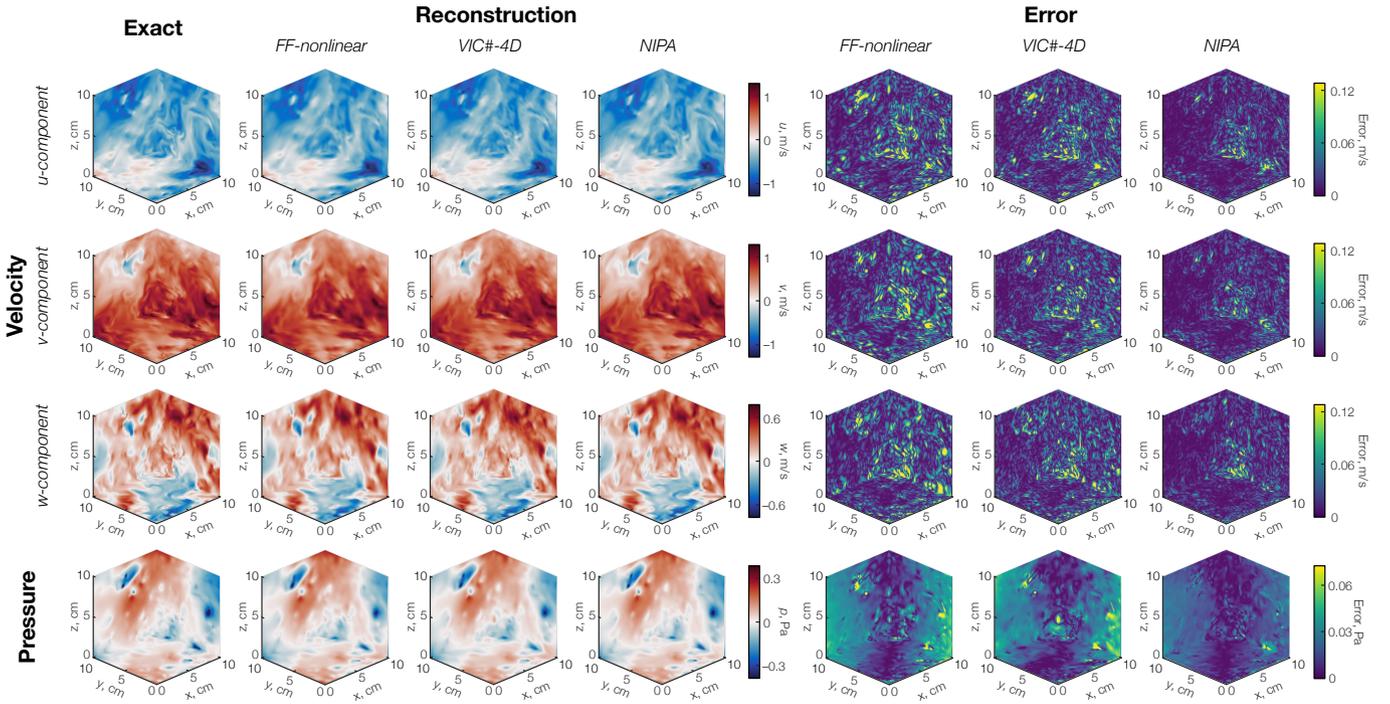
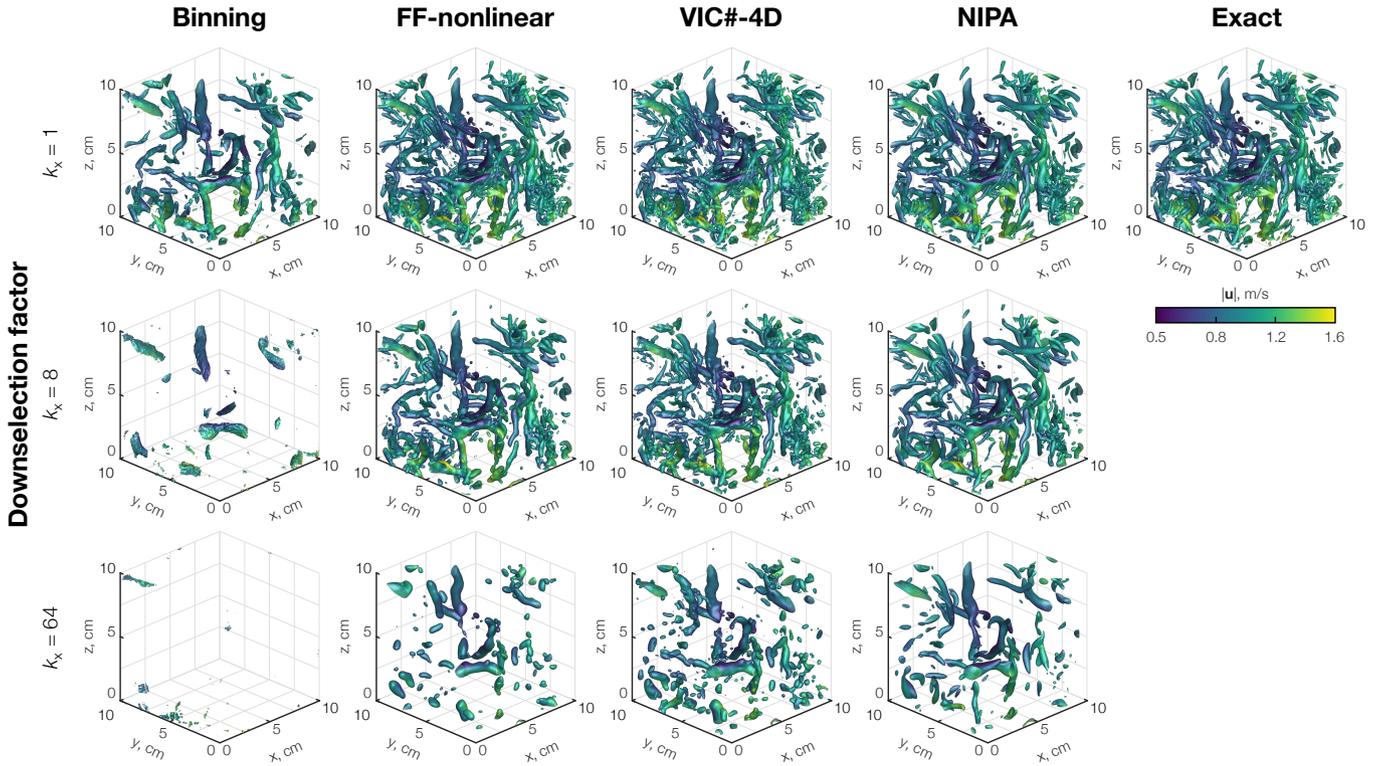


Figure 3. Exact and reconstructed HIT flow fields (left) and absolute error fields (right).

### 4.3. Synthetic TBL results

Next is the synthetic TBL case. While boundary conditions are unknown in the HIT case, a no-slip boundary condition is included at the wall in FlowFit and NIPA for both the synthetic and experimental TBL cases. Specifically, FlowFit includes virtual particles with zero velocity at the boundary surface, while NIPA incorporates an explicit boundary loss to penalize non-zero velocity along the wall. VIC# does not include special treatment of this boundary. NRMSEs of the fluctuating velocity and pressure at the central frame are plotted for all the methods in Fig. 5. Similar trends can be observed as in Sec. 4.3. First, binning errors are generally larger than those of DA methods, except for VIC# at larger time separations. Again, this may be due to the use of sub-optimal hyperparameters in the track filtering routine. Second, all the DA methods are highly sensitive to the particle density, but they are robust to the frame rate and noise variations considered in this work. Note that for the lowest inter-particle distances, both FlowFit and VIC# achieve a lower error than NIPA, which may be due to network's limited expressivity or a sub-optimal network architecture (e.g., the number of layers and neurons, Fourier encoding parameters, etc.). Boosting the expressivity of PINNs is an area of active research and is crucial in flow reconstruction applications. However, in sparse conditions, velocity and pressure NRMSEs are around 30–50% lower for NIPA than for FlowFit and VIC#.

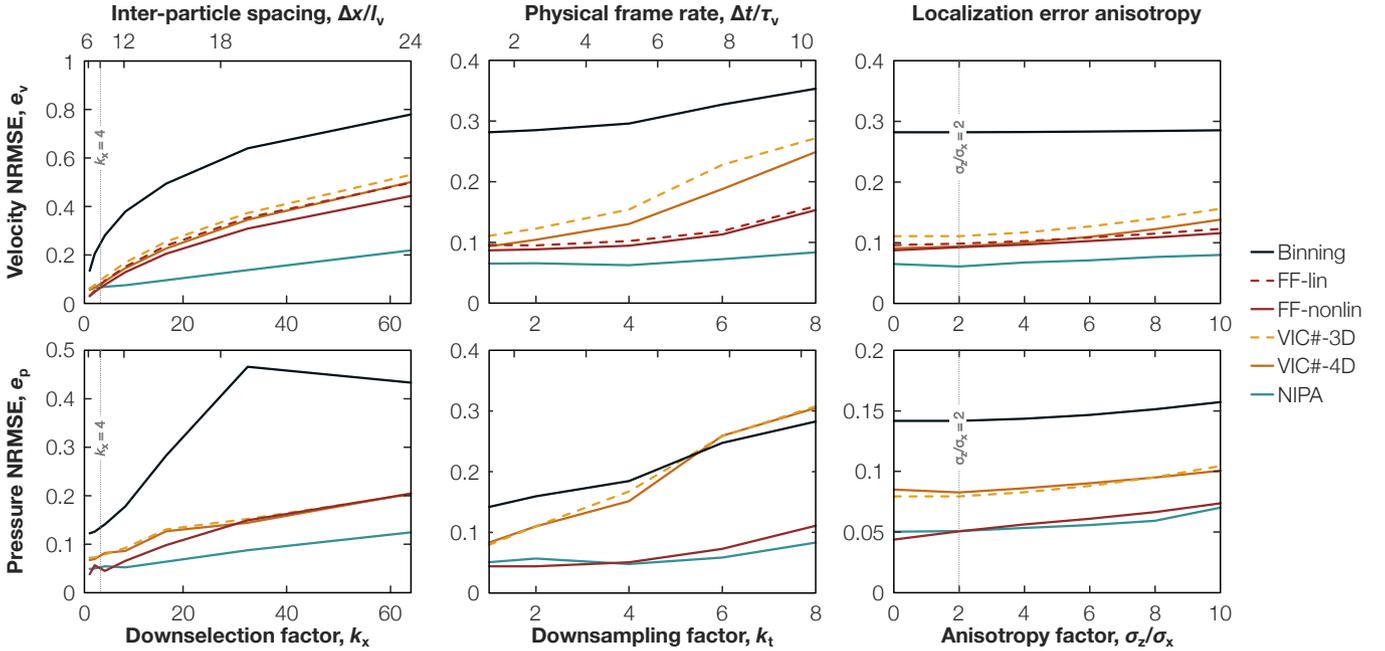
Figure 6 compares cut plots of the TBL velocity and pressure fields at the central frame from the DA methods and the DNS. Once again, a realistic, noisy TBL condition ( $k_x = 4$  and  $\sigma_z/\sigma_x = 2$ ) is selected to demonstrate the performance of DA algorithms. The figure shows three orthogonal planes: a left ( $y = 12$  mm), rear ( $x = 56$  mm), and horizontal ( $z = 0.85$  mm) plane. The DA methods



**Figure 4.** Coherent structures in synthetic HIT flow reconstructed by the binning and DA methods. Structures are visualized using isosurfaces of the  $Q$ -criterion field ( $Q = 3500 \text{ s}^{-2}$ ) colored by the velocity magnitude.

produce high-fidelity estimates of velocity and pressure. NRMSEs of the velocity fluctuations are 9.2%, 9.3%, and 6.1% for FF-nonlinear, VIC#-4D, and NIPA, respectively; pressure errors are 5.1%, 8.2%, and 5.1%. Error fields are shown on the right-side of the figure. Larger errors cluster at the boundaries of the domain for all three DA methods, due to strong in- and outflow through these planes. Both FF-nonlinear and VIC#-4D exhibit a similar spatial distribution of velocity errors, while NIPA achieves slightly lower errors in the horizontal plane ( $z = 0.85 \text{ mm}$ ), especially for the  $u$ -velocity component.

Exact and reconstructed coherent structures based on the  $Q$ -criterion field are plotted in Fig. 7. Binning results and reconstructions from all three methods are compared at  $k_x = 1, 8$ , and  $64$ . Again, binning smooths out fine coherent structures, even at the smallest inter-particle spacing, while the DA methods exhibit near-scale-resolving reconstructions (with NIPA slightly under-resolving some fine structures due to the network's limited expressivity, as mentioned above). At sparser seeding conditions, the DA methods begin to suffer from spatial filtering to different degrees, and binning fails to uncover any meaningful structures. Both FF-nonlinear and VIC#-4D produce an increasing number of high-frequency artifacts that corrupt the true structures, whose shapes are also distorted. By contrast, NIPA better preserves the  $Q$ -criterion field at large downsampling factors, consistent with the observations for the HIT case.



**Figure 5.** Synthetic TBL velocity and pressure NRMSEs at different conditions for the DA and binning methods. Vertical dashed lines indicate the conditions corresponding to Fig. 6.

#### 4.4. Experimental results

We lastly present DA reconstructions of the experimental TBL measurement. Unlike the synthetic cases, where a DNS ground truth is available, we benchmark accuracy in the experimental case using TrackFit results (i.e., filtered particle locations and velocities) using particles in the *test set*. This set contains 20% of the tracked particles, which were unseen by the DA algorithms, i.e., not used to reconstruct the flow, as detailed in Sec. 3.2. Figure 8 shows NRMSEs based on the *test set* velocities at the central frame using different downsampling factors. Again, only the fluctuating component of velocity is used to compute the errors. The ensemble average is obtained by averaging all 51 frames of the NIPA reconstructions based on the full *training set*.

Per Fig. 8, while the accuracy of each method degrades with fewer particles, NIPA produces NRMSEs around 50% lower than those of FlowFit and VIC#. This trend is consistent with trends from the synthetic cases, e.g., as can be seen in the spatial and noise variation panels of Fig. 5 at a realistic ratios ( $\sigma_z/\sigma_x = 2$  to 4). The magnitude of errors in the experimental case are about 40% higher than in the synthetic cases, which is attributable to hardware (image noise, laser intensity fluctuations) and tracking errors. Note that the particle velocity uncertainty from TrackFit is minimal compared to that of the DA-reconstructed velocity field. Past works have reported track velocity errors below 1% of the local root mean square velocity (Schneiders & Sciacchitano, 2017; Janke & Michaelis, 2021), which is significantly less than the DA reconstruction error of 5–10% in Fig. 5. Moreover, the nominal positional uncertainty of STB is less than 0.1 px, which amounts to only  $\sim 0.1l_v$ , i.e., a tenth the diameter of even the finest eddies. These findings support the use of TrackFit data to benchmark the DA reconstructions (Schneiders & Sciacchitano, 2017).

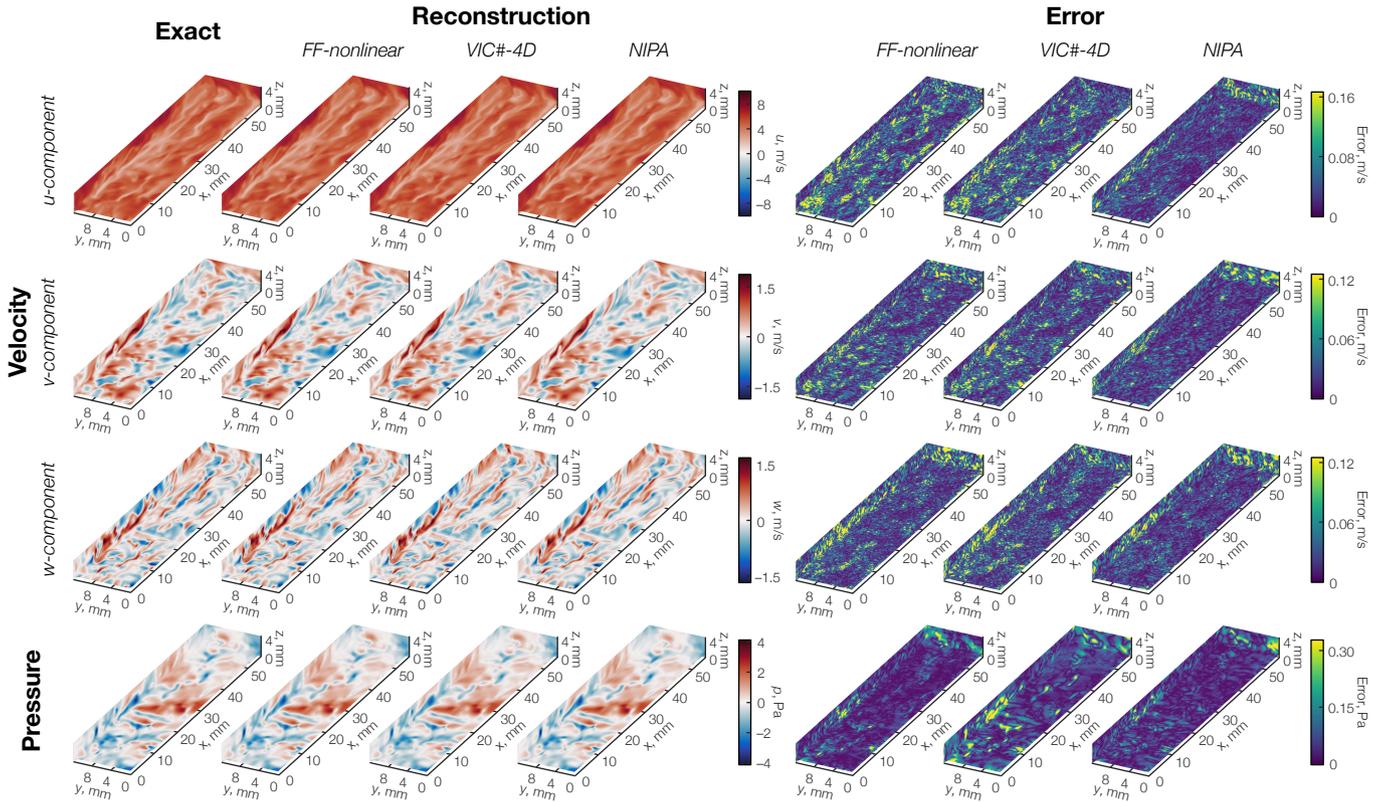
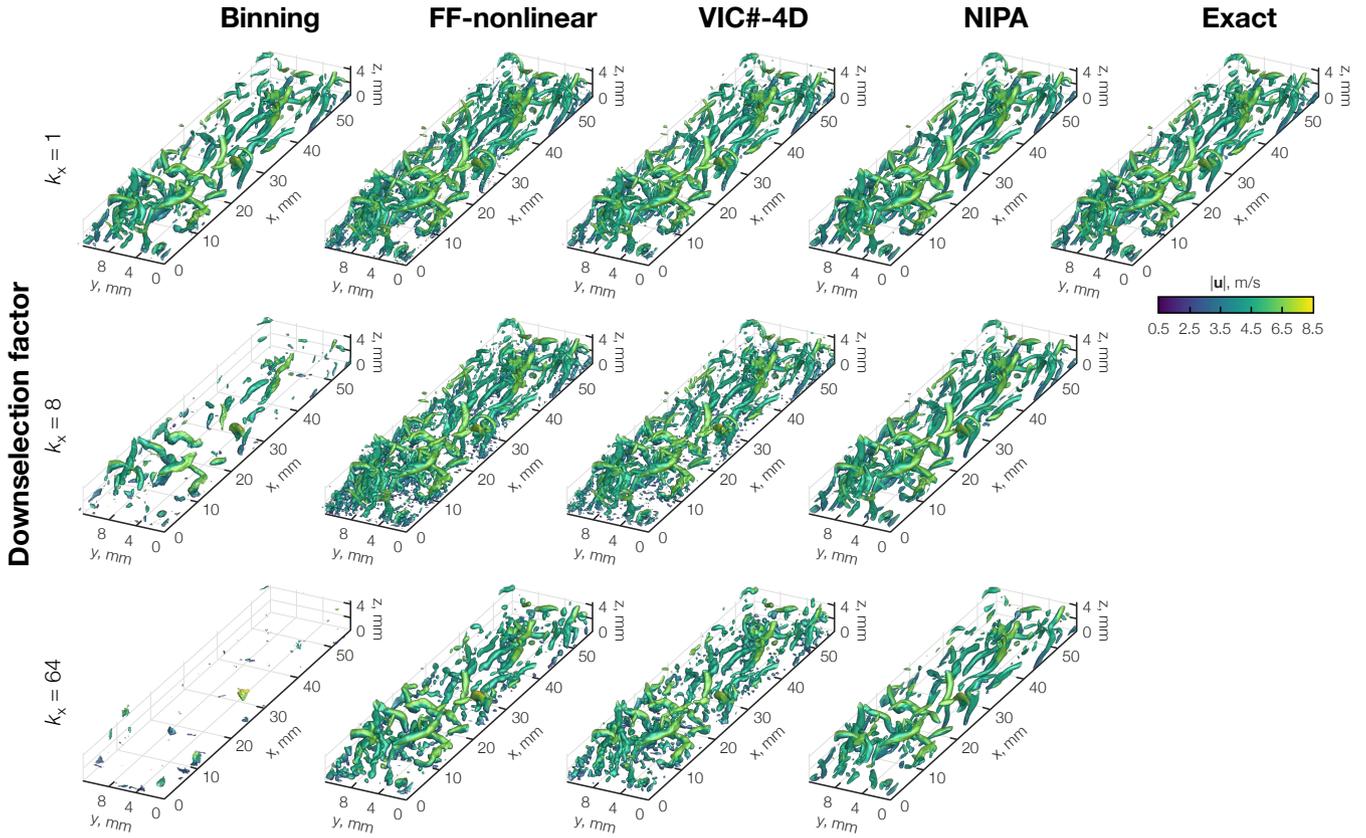


Figure 6. Exact and reconstructed TBL flow fields (left) and absolute error fields (right).

Isosurfaces of the  $Q$ -criterion field are presented in Fig. 9. We plot results for three advanced DA methods and the binning method at two downsampling factors:  $k_x = 1$  (no downsampling of the *training set*) and 8 (significant downsampling). Similar observations can be made as in the synthetic cases: NIPA reconstructs smoother coherent structures that are consistent across downsampling factors, while FF-nonlinear and VIC#-4D are susceptible to some high-frequency errors that obscure the (presumably) true flow features. This corroborates the relative errors in Fig. 8. By contrast, binning is only able to uncover a few bulk structures at the highest particle density; only fragmented vortices are observed in the sparsest case.

In addition to velocity, pressure fluctuations are inferred and compared in Fig. 10. Pressure distributions are shown on the left ( $y = 11.8$  mm) and rear ( $x = 51.8$  mm) planes close the RoI boundaries as well as a half-height plane ( $z = 1.25$  mm). When using the whole *training set*, the DA reconstructions converge to a very similar pressure estimate, which is evidence for the accuracy of these reconstructions.<sup>†</sup> Consistent with our observations about the velocity field, binning yields a low-pass filtered pressure field. As the downsampling factor is increased, binning results are totally corrupted and exhibit wild oscillations due to the large, erroneous velocity gradients shown in Fig. 9. Pressure fields reconstructed by the DA methods also suffer from some erratic fluctuations, though much less so than binning. A salient example is the upstream region of the TBL in Fig. 10,

<sup>†</sup>The volume-averaged pressure variance across the DA reconstructions is  $0.013 \text{ Pa}^2$  and  $0.031 \text{ Pa}^2$  for the  $k_x = 1$  and 8 cases, respectively. This suggests that random errors/uncertainties grow with decreasing seeding density.

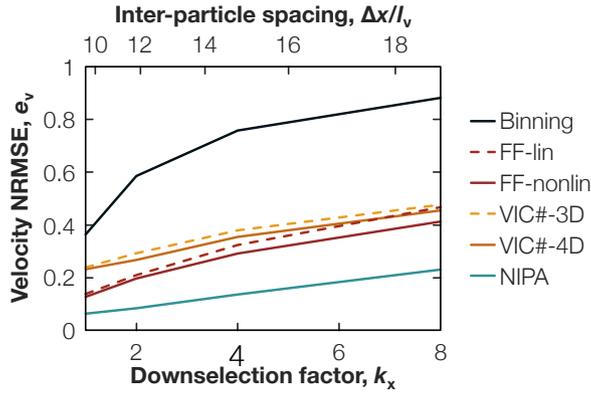


**Figure 7.** Coherent structures in synthetic TBL flow reconstructed by the binning and DA methods. Structures are visualized using isosurfaces of the  $Q$ -criterion field ( $Q = 2.5 \times 10^6 \text{ s}^{-2}$ ) colored by the velocity magnitude.

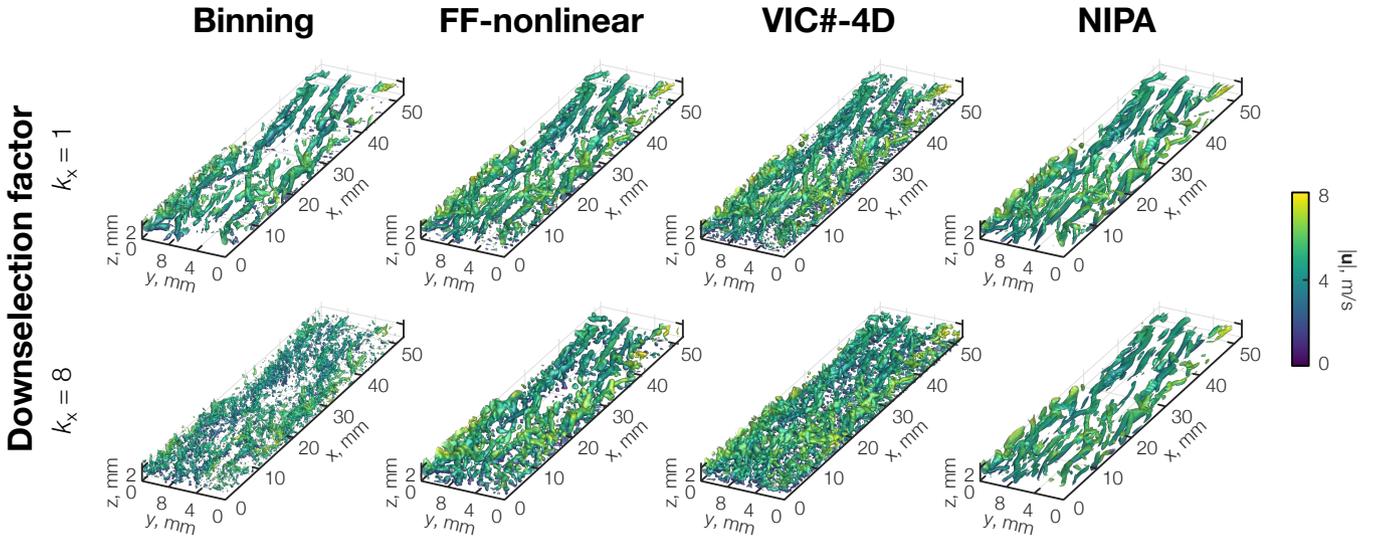
which is highlighted by the red boxes. Over- and underestimation of pressure is qualitatively consistent with the velocity error trends in Fig. 8. VIC# has the largest errors, concentrated at the up- and downstream boundaries; FF-nonlinear has lower errors; and NIPA appears to produce the most accurate estimate. A qualitative error analysis for pressure at  $k_x = 8$  is performed by comparison to the mean high-fidelity estimate (i.e., averaging  $p$  from FF-nonlinear, VIC#-4D, and NIPA in the  $k_x = 1$  case). Pseudo-errors are shown on the bottom row of Fig. 10, where the mean  $k_x = 1$  pressure field (excluding binning) is compared to the  $k_x = 8$  fields;  $e_p$  is 52%, 88%, and 47% for FF-nonlinear, VIC#-4D, and NIPA. Future investigations will consider pressure and pressure error spectra in comparison to DNS results and theoretical predictions.

## 5. Computational cost

It is not possible to provide a controlled comparison of computational cost since the FlowFit algorithms are run on CPUs, while VIC# and NIPA are run on a GPU. Therefore, we briefly discuss the DA algorithms in terms of hardware and wall time to provide a qualitative sense of cost. Both variants of FlowFit are executed on an AMD<sup>®</sup> Ryzen<sup>™</sup> 9 5950X CPU (16 cores), and binning is done on an Intel<sup>®</sup> Xeon<sup>™</sup> w7-3465X (28 cores). VIC# and NIPA are run on NVIDIA<sup>®</sup> GPUs: the former on an RTX<sup>™</sup> 4090 and Tesla<sup>™</sup> V100, the latter on an RTX<sup>™</sup> A6000. Table 2 summarizes the average computation time of each method for a representative case ( $k_x = 4$  and  $\sigma_z/\sigma_x = 2$ ) for both the



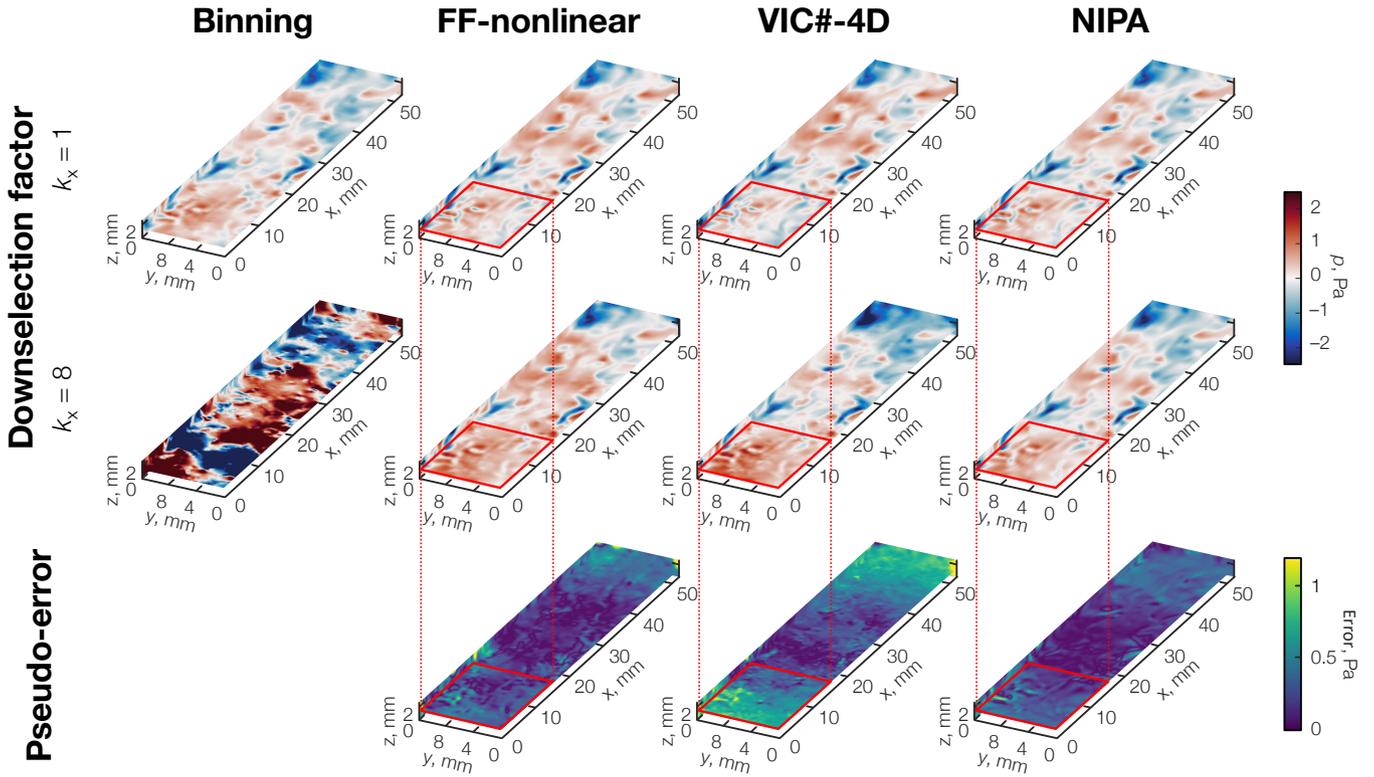
**Figure 8.** Experimental TBL velocity and pressure NRMSEs at different conditions for the DA and binning methods. Errors are computed using the *test set* data from TrackFit, which was not included in the reconstructions.



**Figure 9.** Coherent structures in experimental TBL flow reconstructed by the binning and DA methods. Structures are visualized using isosurfaces of the  $Q$ -criterion field ( $Q = 1.25 \times 10^6 \text{ s}^{-2}$ ) colored by the velocity magnitude.

synthetic flows. Other cases have slightly different costs due to the varying number of tracks, but relative trends are the same. Wall times are normalized by the number of time steps included in the reconstruction, discussed further in Sec. 2.

The methods exhibit quite different runtimes, with binning and FlowFit-linear being the fastest, taking only a couple of seconds on a mainstream CPU. FlowFit-nonlinear takes significantly longer but is still around an order of magnitude faster than the other methods, which are running on GPUs. NIPA takes the longest time of all. This is potentially due to slow convergence of the first-order Adam optimizer, as compared to the quasi-Newton L-BFGS optimizer used by FlowFit and VIC#, as well as a more challenging optimization problem that simultaneously modifies the tracks and flow states.



**Figure 10.** Pressure fields in experimental TBL flow reconstructed by binning and DA methods (top). Pseudo-error fields of the  $k_x = 8$  reconstructions compared to the mean  $k_x = 1$  pressure field (bottom). Highlighted regions indicate large in-flow errors.

**Table 2.** Wall time in seconds of each DA method for the synthetic cases.

	Binning	VIC#-3D	VIC#-4D	FF-linear	FF-nonlinear	NIPA
HIT	13	393	868	10	73	1420
TBL	6	174	585	6	22	850

## 6. Conclusions

This paper presents a comprehensive comparison of three state-of-the-art LPT DA algorithms: FlowFit, VIC#, and NIPA. Synthetic and experimental tests are conducted to systematically assess the performance of these algorithms, which are baselined by a conventional interpolation method (“binning”). Accuracies and costs of the DA methods are investigated in terms of the mean inter-particle distance, temporal resolution, and magnitude of tracking errors. Results show DA’s superiority over binning due to including the governing physics. DA reconstruction errors are found to be highly sensitive to the inter-particle distance but robust to the frame rate and localization errors. Synthetic tests show that the FlowFit and VIC# families produce noisier reconstructions under sparse seeding conditions, while NIPA is comparatively robust to this parameter. The same trend is observed in the experimental TBL case, where NIPA achieves velocity errors around 50% lower than those from FlowFit and VIC#. Computational costs are reported for all three methods. FlowFit is far and away the fastest algorithm, consuming an order of magnitude less time than

VIC# and NIPA, and NIPA takes the longest time. Future work will analyze velocity and pressure error spectra to better understand implicit filtering by the DA algorithms, and we will develop a hybrid DA scheme that can leverage the relative speed and accuracy of the different techniques.

## References

- Agterberg, F. (2014). *Geomathematics: Theoretical Foundations, Applications and Future Developments* (Vol. 18). Springer.
- Agüera, N., Cafiero, G., Astarita, T., & Discetti, S. (2016). Ensemble 3D PTV for high resolution turbulent statistics. *Measurement Science and Technology*, 27(12), 124011.
- Agüí, J. C., & Jimenez, J. (1987). On the performance of particle tracking. *Journal of Fluid Mechanics*, 185, 447–468.
- Cai, S., Gray, C., & Karniadakis, G. E. (2024). Physics-Informed Neural Networks Enhanced Particle Tracking Velocimetry: An Example for Turbulent Jet Flow. *IEEE Transactions on Instrumentation and Measurement*.
- Casa, L., & Krueger, P. (2013). Radial basis function interpolation of unstructured, three-dimensional, volumetric particle tracking velocimetry data. *Measurement Science and Technology*, 24(6), 065304.
- Christiansen, I. P. (1973). Numerical simulation of hydrodynamics by the method of point vortices. *Journal of Computational Physics*, 13(3), 363–379.
- Fahringer, T. W., Lynch, K. P., & Thurow, B. S. (2015). Volumetric particle image velocimetry with a single plenoptic camera. *Measurement Science and Technology*, 26(11), 115201.
- Gesemann, S., Huhn, F., Schanz, D., & Schröder, A. (2016). 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* (Vol. 4, pp. 1–17).
- Giering, R., & Kaminski, T. (1998). Recipes for adjoint code construction. *ACM Transactions on Mathematical Software*, 24(4), 437–474.
- Godbersen, P., Gesemann, S., Schanz, D., & Schröder, A. (2024). FlowFit3: Efficient data assimilation of LPT measurements. In *21st International Symposium on the Application of Laser and Imaging Techniques to Fluid Mechanics, Lisbon, Portugal*.
- Grauer, S. J., & Zhou, K. (2024). *Reconstructing complex flows using neural data assimilation*. Experiments in Fluids Seminar Series.
- Imaichi, K., & Ohmi, K. (1983). Numerical processing of flow-visualization pictures—measurement of two-dimensional vortex flow. *Journal of Fluid Mechanics*, 129, 283–311.

- Janke, T., & Michaelis, D. (2021). Uncertainty quantification for PTV/LPT data and adaptive track filtering. In *14th International Symposium on Particle Image Velocimetry, Virtual* (Vol. 1).
- Jeon, Y. J. (2023). Vorticity time-marching method in fine scale reconstruction (VIC#) for describing 4-D space-time. In *15th International Symposium on Particle Image Velocimetry, San Diego, CA*.
- Jeon, Y. J., Müller, M., & Michaelis, D. (2022). Fine scale reconstruction VIC# by implementing additional constraints and coarse-grid approximation into VIC+. *Experiments in Fluids*, 63(4), 70.
- Li, L., & Pan, Z. (2024). Three-dimensional time-resolved Lagrangian flow field reconstruction based on constrained least squares and stable radial basis function. *Experiments in Fluids*, 65(4), 57.
- Malik, N. A., & Dracos, T. (1995). Interpolation schemes for three-dimensional velocity fields from scattered data using Taylor expansions. *Journal of Computational Physics*, 119(2), 231–243.
- Miozzi, M. (2005). Accuracy of  $C^1$  natural neighbors interpolation of feature tracking data and vorticity measurements. In *6th International Symposium on Particle Image Velocimetry, Pasadena, CA*.
- Molnar, J. P., & Grauer, S. J. (2022). Flow field tomography with uncertainty quantification using a Bayesian physics-informed neural network. *Measurement Science and Technology*, 33(6), 065305.
- Molnar, J. P., Venkatakrishnan, L., Schmidt, B. E., Sipkens, T. A., & Grauer, S. J. (2023). Estimating density, velocity, and pressure fields in supersonic flows using physics-informed BOS. *Experiments in Fluids*, 64(1), 14.
- Perlman, E., Burns, R., Li, Y., & Meneveau, C. (2007). Data exploration of turbulence simulations using a database cluster. In *International Conference for High Performance Computing, Networking, Storage and Analysis, Reno, NV* (pp. 1–11).
- 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.
- Rosales, C., & Meneveau, C. (2005). Linear forcing in numerical simulations of isotropic turbulence: Physical space implementations and convergence properties. *Physics of Fluids*, 17(9), 095106.
- Savitzky, A., & Golay, M. J. E. (1964). Smoothing and differentiation of data by simplified least squares procedures. *Analytical Chemistry*, 36(8), 1627–1639.
- Scarano, F. (2012). Tomographic PIV: principles and practice. *Measurement Science and Technology*, 24(1), 012001.

- Schanz, D., Gesemann, S., & Schröder, A. (2016). Shake-The-Box: Lagrangian particle tracking at high particle image densities. *Experiments in Fluids*, 57, 70.
- Schanz, D., Gesemann, S., Schröder, A., Wieneke, B., & Novara, M. (2012). Non-uniform optical transfer functions in particle imaging: calibration and application to tomographic reconstruction. *Measurement Science and Technology*, 24(2), 024009.
- Schanz, D., Jahn, T., & Schröder, A. (2022). 3D particle position determination and correction at high particle densities. In *20th International Symposium on the Application of Laser and Imaging Techniques to Fluid Mechanics, Lisbon, Portugal* (pp. 1933–1949).
- Schneiders, J. F. G., & Scarano, F. (2016). Dense velocity reconstruction from tomographic PTV with material derivatives. *Experiments in Fluids*, 57, 139.
- Schneiders, J. F. G., & Sciacchitano, A. (2017). Track benchmarking method for uncertainty quantification of particle tracking velocimetry interpolations. *Measurement Science and Technology*, 28(6), 065302.
- Schröder, A., & Schanz, D. (2023). 3D Lagrangian particle tracking in fluid mechanics. *Annual Review of Fluid Mechanics*, 55, 511–540.
- Schröder, A., Schanz, D., Geisler, R., Godbersen, P., Agocs, J., & Simhan, A. R. (2024). Near-wall flow features in ZPG-TBL at various Reynolds numbers using dense 3D Lagrangian Particle Tracking. In *21st International Symposium on the Application of Laser and Imaging Techniques to Fluid Mechanics, Lisbon, Portugal*.
- Schröder, A., Schanz, D., Geisler, R., Novara, M., & Willert, C. (2015). Near-wall turbulence characterization using 4D-PTV “Shake-The-Box”. In *11th International Symposium on Particle Image Velocimetry, Santa Barbara, USA* (pp. 1–14).
- Sciacchitano, A., Leclaire, B., & Schröder, A. (2021). Main results of the first data assimilation challenge. In *14th International Symposium on Particle Image Velocimetry, Virtual* (Vol. 1, pp. 1–13).
- Sciacchitano, A., Leclaire, B., & Schröder, A. (2022). Main results of the analysis of the HOMER Lagrangian Particle Tracking and Data Assimilation database. In *20th International Symposium on the Application of Laser and Imaging Techniques to Fluid Mechanics, Lisbon, Portugal* (pp. 1823–1847).
- Sperotto, P., Pieraccini, S., & Mendez, M. A. (2022). A meshless method to compute pressure fields from image velocimetry. *Measurement Science and Technology*, 33(9), 094005.
- Stüben, K. (2001). An introduction to algebraic multigrid. In U. Trottenberg, C. W. Oosterlee, & A. Schüller (Eds.), *Multigrid* (pp. 413–532). Academic Press.
- Stüer, H., & Blaser, S. (2000). Interpolation of scattered 3D PTV data to a regular grid. *Flow, Turbulence and Combustion*, 64, 215–232.

- Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., ... Ng, R. (2020). Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems*, 33, 7537–7547.
- Toloui, M., Mallery, K., & Hong, J. (2017). Improvements on digital inline holographic PTV for 3D wall-bounded turbulent flow measurements. *Measurement Science and Technology*, 28(4), 044009.
- van Gent, P. L., Michaelis, D., van Oudheusden, B. W., Weiss, P.-É., de Kat, R., Laskari, A., ... Schrijer, F. F. J. (2017). Comparative assessment of pressure field reconstructions from particle image velocimetry measurements and Lagrangian particle tracking. *Experiments in Fluids*, 58, 33.
- Wang, H., Gao, Q., Feng, L., Wei, R., & Wang, J. (2015). Proper orthogonal decomposition based outlier correction for PIV data. *Experiments in Fluids*, 56, 43.
- Wang, S., Sankaran, S., Wang, H., & Perdikaris, P. (2023). An expert's guide to training physics-informed neural networks. *arXiv preprint*, 2308.08468.
- Wang, S., Yu, X., & Perdikaris, P. (2022). When and why PINNs fail to train: A neural tangent kernel perspective. *Journal of Computational Physics*, 449, 110768.
- Wieneke, B. (2008). Volume self-calibration for 3D particle image velocimetry. *Experiments in Fluids*, 45(4), 549–556.
- Zhou, K., & Grauer, S. J. (2023). Flow reconstruction and particle characterization from inertial Lagrangian tracks. *arXiv preprint*, 2311.09076.
- Zhou, K., Li, J., Hong, J., & Grauer, S. J. (2023). Stochastic particle advection velocimetry (SPAV): theory, simulations, and proof-of-concept experiments. *Measurement Science and Technology*, 34(6), 065302.