

Physics-Informed Flow Reconstruction from Sparse 3D Lagrangian Particle Tracking: A Systematic Comparison with FlowFit and Binning

Hyungmin Shin

`hyungmin.shin@dlr.de`

German Aerospace Center (DLR)

Philipp Godbersen

German Aerospace Center (DLR)

Andreas Schröder

German Aerospace Center (DLR)

Research Article

Keywords: PINNs, Data Assimilation, Lagrangian Particle Tracking, Turbulent Flow

Posted Date: October 19th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11270610/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Physics-Informed Flow Reconstruction from Sparse 3D Lagrangian Particle Tracking: A Systematic Comparison with FlowFit and Binning

Hyungmin Shin^{1*}, Philipp Godbersen¹ and Andreas Schröder^{1,2}

^{1*}Institute of Aerodynamics and Flow Technology, German Aerospace Center (DLR), Bunsenstrasse 10, Göttingen, 37073, Germany.

²Institute of Transport Technology, Brandenburg University of Technology Cottbus-Senftenberg, Siemens-von-Halske Ring 15a, Cottbus, 03046, Germany.

*Corresponding author(s). E-mail(s): hyungmin.shin@dlr.de;
Contributing authors: philipp.godbersen@dlr.de;
andreas.schroeder@dlr.de;

Abstract

Reconstructing continuous Eulerian flow fields from sparse three-dimensional Lagrangian Particle Tracking (LPT) data remains challenging, particularly for velocity gradients and pressure. In this study, Physics-Informed Neural Networks (PINNs) are systematically evaluated over a wide range of particle concentrations and compared with polynomial Binning and FlowFit using identical TrackFit-processed LPT data. Two synthetic datasets, homogeneous isotropic turbulence (HIT) and a turbulent boundary layer (TBL), together with an experimental turbulent-boundary-layer dataset, are used to assess velocity, pressure, vorticity, acceleration, incompressibility, and flow-structure reconstruction. PINNs yields the lowest velocity and pressure errors over the investigated sparsity levels and generally shows slower degradation as particle density decreases. Although FlowFit provides exact divergence-free velocity fields and outperforms PINNs in vorticity reconstruction at several densely seeded conditions, PINNs preserves coherent flow structures more robustly over time and under sparse seeding. In the experimental case, PINNs also achieves the lowest validation acceleration error despite not using measured acceleration as a training target. A broader comparison of the PINNs approach used in this study with other flow-reconstruction methods is available in the results of the 2nd LPT and DA challenges.

Keywords: PINNs, Data Assimilation, Lagrangian Particle Tracking, Turbulent Flow

1 Introduction

Lagrangian Particle Tracking (LPT) is a well-established experimental technique for analyzing the trajectories of individual tracer particles, enabling quantitative visualization of fluid flow behavior. Unlike volumetric measurement techniques such as Tomographic Particle Image Velocimetry (Tomo-PIV), which typically estimate velocity fields through cross-correlation of reconstructed particle-intensity distributions over interrogation volumes (Elsinga et al. 2006; Scarano 2013), LPT tracks individual particles as they move through the flow. This approach provides direct access to particle trajectories, velocities, and accelerations with high spatial localization accuracy (Schröder and Schanz 2023). In addition, major advances in LPT, such as the Shake-The-Box (STB) algorithm (Schanz et al. 2016), have enabled time-resolved three-dimensional particle tracking in densely seeded flows, making LPT a powerful tool for investigating complex unsteady flows, including turbulent flows. However, several challenges remain when LPT data are used to analyze and visualize fluid flows. In particular, the estimation of spatial velocity gradients and pressure fields is difficult for LPT data as the measurements are scattered in space and time (Sciacchitano et al. 2025).

Reconstructing a full flow field on a structured Cartesian grid from scattered LPT data, including velocity-gradient components and pressure, is commonly associated with Cartesian grid reconstruction (CGR). Various data assimilation (DA) techniques have been developed for this purpose. A straightforward approach is ensemble or binning-based averaging, in which particle velocity vectors are averaged within small spatial bins (Kasagi and Nishino 1991; Agüera et al. 2016). Although this method is simple and easy to implement, it does not inherently enforce the governing equations and can therefore suffer from noise, loss of small-scale structures, and inaccurate gradients under sparse seeding conditions (Sciacchitano et al. 2025). To incorporate physical constraints into the reconstruction process, FlowFit represents the velocity field using a three-dimensional B-spline basis. The optimal coefficients of the B-spline basis are determined by assimilating measured particle velocities and accelerations together with the continuity and momentum equations (Gesemann et al. 2016). This method has been further improved by enforcing divergence-freeness as a hard constraint, thereby restricting the optimization to a divergence-free subspace of the velocity field (Godbersen et al. 2024). Similarly, Vortex-In-Cell (VIC) based algorithms, such as VIC+ and VIC#, incorporate vorticity dynamics together with measured velocity and acceleration data to improve the physical consistency of the reconstructed flow field (Schneiders et al. 2014; Schneiders and Scarano 2016; Jeon et al. 2022).

With the advent of Physics-Informed Neural Networks (PINNs, Raissi et al. 2019), several studies have demonstrated the potential of neural-network-based DA for LPT and related sparse flow measurements. Wang et al. (2022a) demonstrated PINN-based reconstruction of simulated LPT data generated from direct numerical simulation (DNS) of channel flow and experimental 3D wake flow of a hemisphere. Clark Di Leoni et al. (2023) compared PINNs reconstruction with the conventional DA method Constrained Cost Minimization (CCM) for synthetic DNS channel flow and experimental turbulent shear flow. Cai et al. (2024) showed that PINNs can achieve super-resolution velocity reconstruction in both space and time from sparse measurements obtained in

tomographic PTV experiments of jet flow in water. [Steinfurth et al. \(2024\)](#) showed that the data assimilation of PTV data with PINN is suitable compared to conventional DA method by showing better reconstruction of vortex rings in particular under conditions of severe data sparsity. [Toscano et al. \(2024\)](#) introduced AIVT, a Kolmogorov–Arnold-network-based method for temperature inference in Rayleigh–Bénard convection which demonstrated that neural network based method can infer continuous and instantaneous temperature fields from sparse experimental data. More recently, [Zhou and Grauer \(2026\)](#) proposed joint reconstruction of particle trajectories and Eulerian flow field in multiple types of test cases and showed better reconstruction of flow fields and particle trajectories even when large particle position errors are in the measurement data.

Nevertheless, existing works focus primarily on demonstrating reconstruction feasibility for a given measurement configuration or on comparing PINNs with a limited set of conventional DA methods. In contrast, less attention has been paid to how specific LPT experimental parameters influence the quality of the reconstructed Eulerian field. This point is particularly important because the practical performance of any DA method is not determined only by the reconstruction algorithm itself, but also by the particle concentration, temporal sampling rate, measurement noise, and spatial distribution of the available tracks. Furthermore, while velocity reconstruction has been widely investigated, the preservation of velocity gradients and the accuracy of pressure reconstruction remain more demanding and less systematically understood.

The present study addresses this gap by systematically investigating the relationship between volumetric particle concentration and PINN-based reconstruction performance. Rather than treating PINNs only as a super-resolution tool, we evaluate their ability to recover physically meaningful flow structures, velocity gradients, divergence-free behavior, and pressure fields under controlled variations of particle concentration and corresponding mean inter particle spacing. The study is designed to clarify when PINNs provide clear advantages over conventional DA approaches, when their performance becomes limited by the available LPT information, and how these limitations appear in gradient-sensitive diagnostics. By considering both synthetic LPT datasets generated from numerical reference fields and experimental datasets obtained with the STB method, we compare the reconstruction performance of PINNs with conventional DA algorithms and investigate the influence of realistic experimental constraints on reconstruction quality.

This paper is organized as follows. Section 2 describes the preprocessing procedure used to convert LPT data into a suitable form for data assimilation. Section 3 provides an overview of the data assimilation methods considered in this study, including conventional reconstruction approaches used as baselines. Section 4 presents the PINNs formulation, the governing equations, the loss functions, and the numerical strategies adopted for robust reconstruction from sparse particle data. Section 5 introduces the numerical and experimental test cases, ranging from synthetic LPT datasets generated from reference flow fields to actual STB measurements. Section 6 presents the reconstruction results for different particle concentrations, with emphasis on velocity,

velocity gradients, pressure, and flow-structure accuracy. Finally, Section 7 summarizes the main findings and discusses the implications for future LPT-based flow reconstruction using physics-informed neural networks.

2 LPT data processing

In both synthetic and experimental LPT measurements considered in this study, the STB method is assumed to be used to reconstruct the three-dimensional positions of individual tracer particles. The reconstructed particle positions are associated with unique particle identifiers, which enables particle trajectories to be followed over time. However, the measured particle positions inevitably contain uncertainty in all spatial directions, which is generally anisotropic, with larger errors often occurring in the depth direction (Schanz et al. 2022). Such position noise can deteriorate the accuracy of the reconstructed trajectories and, consequently, the estimation of kinematic quantities such as particle velocity and acceleration.

To mitigate this issue, the TrackFit algorithm is employed as a trajectory reconstruction and smoothing framework (Gesemann et al. 2016; Buchwald et al. 2026). TrackFit represents each particle trajectory using B-spline curves, providing a continuous and smooth approximation of the discrete and noisy position measurements. The spline coefficients are determined through an optimization procedure that balances fidelity to the measured particle positions with a smoothness constraint. In this way, measurement noise can be filtered while the underlying particle motion is preserved over the measurement time. Once the trajectory is represented by B-splines, particle velocity and acceleration can be obtained analytically from the first and second derivatives of the spline functions.

In this study, TrackFit is applied prior to the DA process to obtain smooth particle trajectories and consistent estimates of particle velocity and acceleration. These reconstructed physical quantities, evaluated at the particle positions and corresponding time instances, are subsequently used as inputs to the DA algorithms.

3 Flow reconstruction methods

To evaluate the performance of PINNs in data assimilation, two additional flow reconstruction methods other than PINNs are considered in this study: Binning and FlowFit. All methods use the same TrackFit-processed LPT data described in the previous section, namely particle positions together with the corresponding velocity and, where required, acceleration information.

The reconstruction strategies used in this study are described in the following subsections. The binning-based method is used as a simple grid-based reference approach that reconstructs the Eulerian velocity and acceleration field directly from scattered particle data. FlowFit is considered as a representative physics-based data assimilation method with established performance for LPT-based flow reconstruction. Finally, the PINNs framework is investigated as a mesh-free, physics-constrained alternative that can produce continuous spatio-temporal flow fields.

3.1 Binning

For the binning-based reconstruction, the scattered particle velocity and acceleration data are projected onto a uniform Cartesian grid using a non-overlapping polynomial binning procedure. At each grid point, neighboring particle velocities are approximated by a local second-order polynomial. The fitting is performed using Gaussian distance-based weights and a small ridge regularization term to improve numerical stability. If the number of particles within a local bin is insufficient, additional nearest-neighbor samples are selected using a KD-tree search. To ensure the no-slip boundary condition, virtual particle sets at the boundary with zero velocity are inserted in the data or training set.

The pressure field is reconstructed from the binned velocity field by solving the incompressible pressure Poisson equation on the same Cartesian grid. The source term is evaluated from the reconstructed velocity-gradient tensor according to the incompressible Navier–Stokes equations. Boundary pressure values are estimated by least-squares integration of the pressure gradient along the boundary shell and are imposed as Dirichlet boundary conditions. The resulting Poisson problem is solved using a finite-difference, matrix-free conjugate-gradient solver.

3.2 FlowFit

FlowFit is a physics-based data assimilation method designed to reconstruct Eulerian flow fields from LPT measurements at individual time-steps. The method represents the velocity field using a three-dimensional uniform B-spline basis, which allows the velocity and its spatial derivatives to be evaluated efficiently at arbitrary locations. The B-spline coefficients are optimized such that the reconstructed velocity field fits the measured particle data while satisfying physical constraints derived from the governing equations everywhere in the measurement volume.

In FlowFit3, divergence-freeness is imposed as a hard constraint during optimization. A staggered grid formulation for the B-spline coefficients allows to constrain the solution to zero divergence in a continuous sense, not just at discrete locations. In this formulation, the coefficient fields associated with the individual velocity components are shifted relative to one another by half a grid spacing. The constraint is enforced via a projection into the divergence-free subspace of the B-splines. Instead of simply projecting the solution, this constraint is applied already within the optimizer which prevents it from even exploring non-feasible solutions. This construction restricts the reconstructed velocity field to a divergence-free subspace and removes the need to enforce the continuity equation only as a soft penalty term. Furthermore, staggered grids places velocities right between pressure nodes, which forces a direct mathematical link. The B-spline coefficients are optimized using a limited-memory Broyden–Fletcher–Goldfarb–Shanno (L-BFGS) algorithm, which provides stable and efficient convergence.

FlowFit can be operated in two modes, commonly referred to as linear and non-linear modes. In the linear mode, the reconstruction enforces divergence-freeness but does not include the momentum equation in the optimization. In the nonlinear mode,

the divergence-free constraint is retained, while the momentum equation and measured particle acceleration are additionally included as soft constraints. This enables the reconstruction of the pressure field through the momentum balance and a pressure Poisson formulation.

In the present study, the nonlinear mode of FlowFit is employed in order to evaluate not only the reconstructed velocity field but also the pressure field. The B-spline grid resolution is adjusted according to the particle concentration so that a consistent particle-per-cell ratio is maintained across the different sparsity levels considered in this study. Additionally, to enforce the no-slip boundary condition, as Binning, virtual particles are generated and inserted at the wall boundary in the training set.

3.3 PINNs

Among neural-network-based machine learning approaches, PINNs provide a framework for approximating unknown solution fields while embedding the governing physical laws into the training process. In this study, the PINN is used as a continuous function approximator that maps spatio-temporal coordinates to the corresponding flow variables. For the LPT-based reconstruction problem, the input consists of the particle position and time, while the network output contains the velocity components and pressure. Since the neural network represents the flow variables as differentiable functions of space and time, the required spatial and temporal derivatives can be obtained directly by automatic differentiation without introducing an additional numerical discretization.

Figure 1 shows the schematic of the PINN method used in this study. The PINN is trained by minimizing a loss function composed of data and physics-based residual terms. Since particle velocities are available from the TrackFit-processed LPT data, the discrepancy between the measured particle velocity and the predicted velocity is included as the data loss,

$$L_d = \frac{1}{N_d} \sum_{j=1}^{N_d} |\hat{\mathbf{u}}_j - \mathbf{u}_j|^2 \quad (1)$$

where $\mathbf{u}$, $\hat{\mathbf{u}}$, and N_d denote the measured velocity, predicted velocity, and batch size of data points, respectively.

In addition to fitting the measured particle velocities, the PINNs enforce physical constraints such as the divergence-free condition and the Navier–Stokes momentum equations. The required gradients of velocity and pressure are computed by automatic differentiation of the neural-network outputs with respect to the input coordinates. Since the governing equations can be written explicitly in residual form, their target values are zero. Therefore, the discrepancy between the calculated residuals and zero defines the physics-based equation loss terms of the PINNs.

$$L_i = \frac{1}{N_p} \sum_{j=1}^{N_p} |\nabla \cdot \hat{\mathbf{u}}_j|^2 \quad (2)$$




Fig. 1 Schematic of PINNs. The inputs consist of time t and spatial coordinates $\mathbf{x}$, and the outputs are velocity $\mathbf{u}$ and pressure p . The neural network is trained using the continuity equation and the incompressible Navier–Stokes equations as physical constraints, in combination with the provided LPT dataset.

$$L_m = \frac{1}{N_p} \sum_{j=1}^{N_p} \left| \frac{\partial \hat{\mathbf{u}}_j}{\partial t} + (\hat{\mathbf{u}}_j \cdot \nabla) \hat{\mathbf{u}}_j + \frac{\nabla \hat{p}_j}{\rho} - \nu \nabla^2 \hat{\mathbf{u}}_j \right|^2 \quad (3)$$

where L_i , L_m , N_p , p , ρ , and ν denote the divergence loss, the Navier–Stokes momentum loss, the number of collocation points in a batch, pressure, fluid density, and kinematic viscosity, respectively.

The total loss function is constructed by combining the data loss and the physics-based loss terms. Since the individual loss terms can have different magnitudes, weighting factors are introduced to balance their relative contributions during training. In this study, these weighting factors are selected manually by considering the characteristic scales of the corresponding loss terms. The final loss function is then written as follows.

$$L = \lambda_d L_d + \lambda_i L_i + \lambda_m L_m \quad (4)$$

where λ_d , λ_i , λ_m are weighting factors that control the influence of each term.

4 PINNs design

The accuracy and robustness of a PINN-based reconstruction are strongly influenced by the training strategies and hyperparameters used during optimization. This is especially important for turbulent flow reconstruction from LPT data, where the measurements are sparse and irregularly distributed, while the target flow contains multi-scale structures and strong local velocity gradients. Under these conditions, a standard PINNs formulation may suffer from spectral bias, insufficient enforcement of the governing equations in localized regions, and slow or unstable convergence caused by the multi-component loss function (Krishnapriyan et al. 2021; Wang et al. 2021a, 2022b). These issues can directly affect the reconstruction of small-scale flow structures, velocity gradients, divergence-freeness, and pressure fields.

To address these challenges, several training strategies are incorporated into the present PINNs framework. In this section, we describe the network architecture and

the training procedures adopted in this study, with emphasis on the strategies that were found to improve reconstruction accuracy, physical consistency, and convergence stability.

4.1 Overall architecture

Nowadays, various neural-network architectures have been proposed to improve expressivity and computational efficiency. However, some architectures can be highly sensitive to hyperparameter choices, which may lead to unstable training or prevent convergence of the loss function (Ansar and Ashraf 2025). In this study, we employ a fully connected Multi-Layer Perceptron (MLP), which has been widely used and shown to be effective in physics-informed data assimilation problems. An MLP is composed of a sequence of affine transformations and nonlinear activation functions. The affine transformation at the k^{th} layer is defined as

$$N_k(\mathbf{z}^{k-1}) = \mathbf{W}^k \mathbf{z}^{k-1} + \mathbf{b}^k \quad (5)$$

where $\mathbf{W}^k$ and $\mathbf{b}^k$ are weights and biases of the k^{th} layer. The network output is obtained by the sequential composition of these affine transformations and activation functions,

$$\hat{\mathbf{q}}(\mathbf{z}; \theta) = (N_k \circ \sigma \circ N_{k-1} \circ \cdots \circ \sigma \circ N_1)(\mathbf{z}) \quad (6)$$

here $\circ$ is the composition operator, and $\theta = \{\mathbf{W}^k, \mathbf{b}^k\}_{k=1}^D$ represents all trainable parameters. The input vector $\mathbf{z}$ consists of the spatial coordinates and time, while $\hat{\mathbf{q}}$ denotes the predicted flow variables. In the present study, the network output consists of the three velocity components and pressure, namely $\hat{\mathbf{q}} = (\hat{u}, \hat{v}, \hat{w}, \hat{p})$. We use an MLP with 10 hidden layers, each containing 400 neurons, and the hyperbolic tangent function is used as the activation function. The weights are initialized using Glorot uniform initialization, while the biases are initialized to zero (Glorot and Bengio 2010).

Input scaling plays an important role in the stable training of PINNs, particularly when the input coordinates have different characteristic length or time scales. When raw coordinates are used directly, input variables with larger numerical ranges can dominate the pre-activation values in the hidden layers, which may deteriorate the conditioning of the optimization problem (LeCun et al. 1998). Furthermore, for bounded activation functions such as tanh, excessively large pre-activation values can result in activation saturation and weak gradient propagation. Therefore, each input coordinate is normalized independently to the range $[-1, 1]$ using the min-max transformation,

$$x_{\text{rescale}} = 2 \frac{x - x_{\min}}{x_{\max} - x_{\min}} - 1 \quad (7)$$

where $x_{\min}$, $x_{\max}$ are minimum and maximum values of x in the computational domain, respectively.

Weight normalization is also employed to stabilize the training process. Since the magnitude of the network weights directly affects the scale of hidden activations and their gradients, uncontrolled growth or shrinkage of the weights can lead to exploding or vanishing signals across layers. To mitigate this issue, each weight vector is

reparametrized by separating its magnitude and direction as

$$\mathbf{W} = \frac{g}{\|\mathbf{v}\|} \mathbf{v} \quad (8)$$

where $\mathbf{v}$ is an unconstrained weight matrix, g is a fixed scalar gain factor, and $\|\mathbf{v}\|$ denotes the norm of $\mathbf{v}$ (Salimans and Kingma 2016).

4.2 Fourier feature mapping

Standard MLP-based neural networks are known to exhibit spectral bias, meaning that low-frequency components of the target function tend to be learned more easily than high-frequency components. In turbulent flow reconstruction, this behavior can lead to overly smooth velocity and pressure fields, especially when the available LPT data are sparse and irregularly distributed. As a result, small-scale flow structures and sharp velocity gradients, which are particularly important in near-wall regions, may not be sufficiently resolved.

To mitigate this limitation, a Fourier feature mapping is introduced before the first hidden layer of the MLP. Instead of feeding the normalized input coordinates directly into the network, the input vector is first mapped to a higher-dimensional feature space using sinusoidal functions with fixed random frequencies (Cao et al. 2019; Wang et al. 2021b).

$$\gamma(\mathbf{z}^0) = [\mathbf{z}^0, \sin(2\pi\mathbf{w}_1 \cdot \mathbf{z}^0), \cos(2\pi\mathbf{w}_1 \cdot \mathbf{z}^0), \dots, \sin(2\pi\mathbf{w}_m \cdot \mathbf{z}^0), \cos(2\pi\mathbf{w}_m \cdot \mathbf{z}^0)]^\top \quad (9)$$

where $\mathbf{z}^0$ denotes the normalized input vector consisting of space and time coordinates, $\mathbf{w}_i$ is the i -th randomly sampled frequency vector, and m is number of Fourier modes. The frequency vectors are sampled before training and kept fixed throughout the optimization. The mapped feature vector $\gamma(\mathbf{z}^0)$ is then used as the input to the MLP. By projecting the input coordinates into a sinusoidal feature space, the network can represent higher-frequency variations more effectively, which helps improve the reconstruction of small-scale flow structures and localized gradients.

4.3 Collocation points

In PINNs, the governing equations are enforced by evaluating the PDE residuals at a finite set of collocation points distributed throughout the spatio-temporal domain. Consequently, the number and distribution of these points directly influence how effectively the physical constraints are enforced during training. A commonly used approach is uniform random sampling; however, this strategy treats all regions of the domain equally, regardless of the local magnitude of the equation residual. As training progresses, computational effort may therefore be unnecessarily allocated to regions where the governing equations are already well satisfied, while regions exhibiting larger residuals remain insufficiently sampled (Nabian et al. 2021). To improve the efficiency of physics enforcement, we employ an adaptive sampling strategy based on the Residual-based Adaptive Distribution (RAD) method (Wu et al. 2023). Rather

than replacing the original uniformly distributed collocation points, additional collocation points are generated according to the residual-based probability distribution and appended to the existing collocation-point set.

The corresponding probability density function (PDF) is defined as

$$p(\mathbf{x}) \propto \frac{\varepsilon^k(\mathbf{x})}{\mathbb{E}[\varepsilon^k(\mathbf{x})]} + c \quad (10)$$

where $\varepsilon(\mathbf{x})$ is the PDE residual, k and c are hyperparameters, and $\mathbb{E}[\varepsilon^k(\mathbf{x})]$ represents the mean value of the k^{th} power of the residual over the sampling domain. In practice, this expectation is approximated by averaging 10,000 uniformly sampled candidate points.

In the present study, the network is initially trained using a fixed set of uniformly distributed collocation points. During training, additional collocation points are periodically generated according to the RAD probability distribution, such that locations exhibiting larger equation errors have a higher probability of being sampled. These RAD-selected points are accumulated and subsequently added to the original uniformly distributed collocation-point set. Figure 2 illustrates the spatial distribution of the equation loss at the RAD-selected points at different stages of training. Relatively large equation errors are initially observed at several locations, particularly near the domain boundaries. As training progresses and RAD-selected points are incorporated into the collocation set, the equation loss at these locations decreases substantially.

Figure 3 quantitatively compares the evolution of the mean equation loss evaluated at the RAD-selected points for PINNs trained with and without the RAD augmentation. Before RAD is activated at 2×10^4 training iterations, both models exhibit nearly identical equation losses at these locations. After the RAD-selected points begin to be incorporated into the collocation set, the equation loss decreases more rapidly for the PINN trained with RAD than for the PINN trained using only the uniformly distributed collocation points.

4.4 Optimization

Another key challenge in PINN training arises from the simultaneous optimization of multiple loss terms. When the data loss dominates the optimization, the network may fit the particle measurements well while insufficiently enforcing the governing equations. Conversely, excessive dominance of the physics losses may drive the reconstruction away from the observations or promote overly smooth solutions. To improve the optimization under such competing objectives, we employ Shampoo with Adam in the Preconditioner’s Eigenbasis (SOAP, Vyas et al. 2025; Wang et al. 2026)). SOAP combines Adam-type adaptive updates with curvature-informed preconditioning based on second-order gradient statistics. For a layer with weight matrix $\mathbf{W}_t \in \mathbb{R}^{m \times n}$ and gradient $\mathbf{G}_t \in \mathbb{R}^{m \times n}$, the left and right preconditioning matrices are estimated as

$$\mathbf{L}_t = \beta_2 \mathbf{L}_{t-1} + (1 - \beta_2) \mathbf{G}_t \mathbf{G}_t^\top, \quad \mathbf{R}_t = \beta_2 \mathbf{R}_{t-1} + (1 - \beta_2) \mathbf{G}_t^\top \mathbf{G}_t, \quad (11)$$




Fig. 2 Spatial distribution of the equation loss evaluated at 3,000 RAD-selected collocation points during training. The collocation points are projected onto the x - y plane, and the color indicates the corresponding equation loss.

where β_2 is the exponential moving-average coefficient. These matrices are decomposed as

$$\mathbf{L}_t = \mathbf{Q}_L \mathbf{\Lambda}_L \mathbf{Q}_L^\top, \quad \mathbf{R}_t = \mathbf{Q}_R \mathbf{\Lambda}_R \mathbf{Q}_R^\top. \quad (12)$$

The gradient is projected into the eigenspace:

$$\tilde{\mathbf{G}}_t = \mathbf{Q}_L^\top \mathbf{G}_t \mathbf{Q}_R. \quad (13)$$

Adam updates are then performed in this rotated space:

$$\tilde{\mathbf{W}}_{t+1} = \tilde{\mathbf{W}}_t - \eta \text{Adam}(\tilde{\mathbf{G}}_t), \quad (14)$$

$$\mathbf{W}_{t+1} = \mathbf{Q}_L \tilde{\mathbf{W}}_{t+1} \mathbf{Q}_R^\top. \quad (15)$$

This procedure provides a geometry-aware optimization direction that can improve convergence in the highly coupled loss landscape encountered in PINN training.

To further examine the optimization behavior, we calculated the alignment scores of SOAP and compared them with those of Adam. The alignment score of a given set of vectors is defined as

$$\mathcal{A}(\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n) = 2 \left\| \frac{1}{n} \sum_{i=1}^n \frac{\mathbf{v}_i}{\|\mathbf{v}_i\|} \right\|^2 - 1, \quad (16)$$




Fig. 3 Evolution of the mean equation loss evaluated at the RAD-selected collocation points for PINNs trained with and without RAD. RAD is activated at 2×10^4 training iterations, as indicated by the vertical dashed line.

where $\mathbf{v}_n$ denotes a gradient vector with respect to the trainable model parameters. Since this score ranges from $[-1, 1]$, values close to 1 indicate strongly aligned gradient directions. The intra-step and inter-step alignment scores are defined as

$$\mathcal{A}_{\text{intra}}^k = \mathcal{A}(\mathbf{g}_1^k, \mathbf{g}_2^k, \dots, \mathbf{g}_n^k), \quad (17)$$

$$\mathcal{A}_{\text{inter}}^k = \mathcal{A}(\mathbf{g}^{k-1}, \mathbf{g}^k), \quad (18)$$

where $\mathbf{g}_n^k$ denotes the gradient of the n th individual loss term and $\mathbf{g}^k$ denotes the gradient of the total loss at training step k .

Figure 4 (a) and (b) present the intra-step and inter-step alignment scores of the Adam and SOAP optimizers obtained from one of our test cases—synthetic HIT. The data were collected every 10 training steps, and Gaussian-smoothed curves are shown to illustrate the trends in the alignment scores. The present PINN optimization exhibits different behaviors for the two alignment measures. For the intra-step alignment, SOAP initially produces higher scores than Adam, but the scores become lower during the later stages of training. On the other hand, SOAP maintains substantially higher inter-step alignment throughout most of the training process. Despite the reduced intra-step alignment at later stages, the histories of both the data and equation losses in Figure 4 (c) and (d) show faster convergence with SOAP than with Adam.




Fig. 4 Comparison of gradient alignment and loss convergence for Adam and SOAP in the synthetic HIT case. (a) Intra-step alignment score and (b) inter-step alignment score during training. Alignment scores were evaluated every 10 training steps; the transparent curves show the sampled values, while the solid curves represent Gaussian-smoothed trends. (c) Data loss and (d) equation loss as functions of training time for Adam and SOAP.

5 Data description

Three test cases are considered in this study in order to evaluate the reconstruction performance of the different data assimilation methods under both controlled and realistic LPT conditions. The first two cases are synthetic LPT datasets generated from DNS reference fields, while the third case is based on experimental STB measurements. The synthetic cases provide full reference velocity and pressure fields, allowing quantitative error evaluation of both velocity and pressure reconstruction. In contrast, the experimental case introduces realistic measurement uncertainty, particle-distribution inhomogeneity, and the absence of ground-truth pressure information. The methodology for error estimation is based on track benchmarking the velocity vector data given at omitted particle positions of the input or training data (Schneiders and Sciacchitano 2017). The main physical and measurement parameters of all test cases are summarized in Table 1.

5.1 Synthetic particle tracks in DNS of HIT

The first synthetic case is forced homogeneous isotropic turbulence (HIT) at a Taylor-microscale Reynolds number of $Re_\lambda = 433$. This case is selected as a canonical

turbulent flow with three-dimensional, multi-scale structures and without mean shear or wall effects. The reference DNS field is obtained from the Johns Hopkins Turbulence Database (JHTDB). A central sub-volume is extracted from the full periodic domain, and synthetic tracer particles are advected through this volume to generate LPT-like trajectories. The domain size, Kolmogorov time scale (τ_η), Kolmogorov length scale (l_η), and other important parameters are listed in Table 1.

5.2 Synthetic particle tracks in DNS of TBL

The second synthetic case is a wall-bounded turbulent flow at a friction Reynolds number of $Re_\tau = 1000$. Compared with the HIT case, this dataset contains strong mean shear, anisotropic turbulent structures, and large velocity gradients in the near-wall region. It is therefore used to evaluate the performance of the reconstruction methods under conditions that are more demanding for gradient estimation and flow-structure preservation. As in the HIT case, the reference DNS field is obtained from JHTDB which enables quantitative evaluation of both velocity and pressure reconstruction errors. For the pressure-error evaluation, the imposed mean pressure gradient used in the DNS was taken into account before comparing the reconstructed and reference pressure fields. The extracted domain size, viscous time scale (τ_ν), viscous length scale (l_ν), and relevant flow parameters are given in Table 1.

5.3 Experimental particle tracks of TBL

The third case is an experimental turbulent boundary layer (TBL) dataset measured using the STB method in the one-meter wind tunnel at DLR Göttingen (Schröder et al. 2024). The test section was 3,000 mm long with a cross-section of 740 mm $\times$ 1,000 mm. Air entered the test section at a free-stream velocity of $u_\infty = 10$ m/s and passed a trip device that initiated the development of a turbulent boundary layer along the lower wall. Corresponding $Re_\tau = 995$ and measurements were taken through a window located 2.08 m downstream of the trip.

Di-Ethyl-Hexyl-Sebacat (DEHS) particles with a mean diameter of approximately $1.5 \mu\text{m}$ were generated using a Laskin-nozzle seeder and homogeneously distributed in the tunnel. A dual-head nanosecond laser (Photonics Industries DM200) illuminated the particles from the downstream direction as shown in Figure 5 (a). Images were captured through a bottom window using five high-speed cameras (four Phantom v2640, one v1840) with Scheimpflug adapters as depicted in Figure 5 (b), operating at 23.9 kHz and a spatial resolution of 34.185 px/mm. The measurement volume was $60 \times 2 \times 15 \text{ mm}^3$ in the streamwise, wall-normal, and spanwise directions, respectively. Camera calibration was performed using a LaVision Type 11 3D calibration plate, followed by volume self-calibration and optical transfer function refinement. Particle image densities ranged from 0.06 to 0.07 particle per pixel (ppp). The STB algorithm tracked 50,000–60,000 particles over up to 140,000 consecutive images.

Unlike the synthetic cases, the experimental dataset contains measurement noise, calibration uncertainty, possible track losses, and non-uniform particle distributions.




Fig. 5 Experimental setup in the 1-meter wind tunnel at DLR Göttingen. (a) Illumination of the measurement volume near the wall using two Photonic Industries DM200 lasers, with indications of the flow direction and imaged region. (b) Five high-speed cameras—four Phantom v2640 and one Phantom v1840—record particle light scattering within the measurement volume through an optical window.

Moreover, no complete Eulerian ground-truth field is available. Therefore, the experimental case is not used for pressure-error evaluation, but rather for assessing velocity reconstruction consistency and generalization to withheld particle tracks.

5.4 LPT dataset

For the synthetic datasets particles are treated as ideal Lagrangian tracers. Synthetic tracks are generated by uniformly seeding neutrally buoyant particles in an extended computational domain and advecting them using a fourth-order Runge–Kutta scheme. The extended domain is used to reduce boundary effects near the interrogation volume. After trajectory generation, only particles located inside the interrogation volume are retained for the final LPT dataset. For all cases, 51 consecutive time steps are used for data assimilation. The HIT dataset contains approximately 50,000 particles per frame, while the synthetic TBL dataset contains approximately 280,000 particles per frame. Assuming a 1-MP high-speed camera, these particle numbers correspond to equivalent particle image densities of approximately 0.1 ppp for the HIT case and 0.28 ppp for the synthetic TBL case, respectively. To simulate particle tracking experiment situation, zero-mean Gaussian errors in x and y direction of 0.01 mm and 0.0033 mm have applied for the HIT and TBL cases, respectively. To the z direction which is normally considered as the image normal direction have twice the position error than the other direction. For the experimental TBL case, 51 consecutive time steps are also extracted, with an average of approximately 54,000 tracked particles per frame.

The extracted particle trajectories are further processed using the B-spline-based TrackFit algorithm. TrackFit provides smoothed particle trajectories and analytically evaluated first and second temporal derivatives. Therefore, the particle positions, velocities, and, where required, accelerations are used as inputs to the reconstruction methods. Using the same TrackFit-processed data for all methods ensures that differences in reconstruction performance arise from the data assimilation algorithms

Table 1 Detailed flow parameters of the test cases.

Case	Re_λ or Re_τ	Domain, mm ³	τ_η or τ_ν , ms	l_η or l_ν , mm	ppp range
HIT	433	$100 \times 100 \times 100$	8.20	0.350	0.100–0.001
TBL _{syn}	1000	$56 \times 4.25 \times 12$	0.09	0.036	0.260–0.004
TBL _{exp}	995	$60 \times 2 \times 15$	0.09	0.037	0.042–0.001

themselves rather than from differences in trajectory preprocessing. In the experimental TBL case, due to the absent of ground truth, 80% of the tracked particles are extracted as training set and the rest of the tracks are considered as validation set for error calculation and excluded in reconstruction process (Schneiders and Sciacchitano 2017).

To investigate the influence of particle sparsity, the available training particles are systematically down-sampled while keeping the reference flow field unchanged. The down-sampling factor is defined as $k_n = 2^n$. Thus, only $1/k_n$ of the original training particles are used for reconstruction at each sparsity level. The validation or reference data are kept unchanged for error evaluation.

5.5 Error metric

For the quantitative comparison, the reconstruction accuracy is evaluated using a normalized root-mean-square error,

$$e_\phi = \left(\frac{\langle \|\phi_{\text{pred}} - \phi_{\text{exact}}\|_2^2 \rangle}{\langle \|\phi_{\text{exact}}\|_2^2 \rangle} \right)^{1/2}, \quad \langle \phi \rangle = \frac{1}{|\Omega|} \int_\Omega \phi(\mathbf{x}) \, d\Omega, \quad (19)$$

where ϕ_{pred} and ϕ_{exact} denote the fluctuating components of the predicted and reference fields, respectively, obtained by subtracting the spatial mean from the corresponding field. For vector-valued quantities, such as velocity, the L_2 norm is evaluated over all components. For scalar quantities, such as pressure, the expression reduces to the scalar difference. The divergence residual is evaluated using the following equation:

$$e_{\text{div}} = \langle (\nabla \cdot \mathbf{u}_{\text{pred}})^2 \rangle^{1/2}, \quad (20)$$

where $\mathbf{u}_{\text{pred}}$ denotes the velocity vector of the reconstructed flow field.

In the synthetic datasets, where full reference fields are available, the errors are evaluated for both the fluctuating velocity and pressure. For the experimental TBL case, the reconstructed velocity field is evaluated at the withheld validation particle positions and compared with the corresponding TrackFit-estimated particle velocities.

6 Results

6.1 Evaluation strategy

Unlike the other reconstruction schemes, the PINN is formulated as a time-resolved method that exploits the temporal gradients of the flow field. For a finite temporal reconstruction window, when causality is not explicitly imposed, the reconstruction error tends to increase slightly near both temporal boundaries. This behavior can be attributed to the reduced temporal support at the first and last snapshots, where the reconstruction cannot fully benefit from neighboring information in both temporal directions. The middle timestep of the temporal window where the error remains nearly plateaued is selected for the instantaneous field comparison to reduce the influence of temporal-boundary effects and to provide a representative comparison under balanced temporal support.

6.2 Synthetic HIT results

Figure 6 (a) – (d) show the fluctuation velocity, pressure, vorticity, and normalized divergence residual, respectively. The fluctuation velocity error increases with increasing k_n for all three reconstruction methods. PINN consistently yields the lowest fluctuation velocity NRMSE over the entire range, followed by FlowFit and Binning. A similar trend is observed for pressure. PINN maintains a substantially lower pressure error than the other methods. Although, in the present implementation, PINN incorporates the known forcing term through the momentum equation residual, whereas FlowFit and Binning do not explicitly impose the same momentum balance during the reconstruction, the impact of the forcing term is negligible in pressure reconstruction since it canceled out during Poisson equation generation.

The vorticity error is considerably larger than the velocity error for all methods and increases rapidly as the particle concentration decreases. This reflects the greater sensitivity of velocity gradient quantities to spatial smoothing. Nevertheless, PINN consistently provides the lowest vorticity NRMSE. The increasing vorticity error is also consistent with the progressive loss of small-scale structures observed in the Q-criterion visualizations in Figure 9.

The normalized divergence residual in Figure 6 (d) should be interpreted as a measure of physical consistency rather than reconstruction accuracy. FlowFit produces residuals close to machine precision, consistent with its divergence-free projection on the staggered grid. PINN yields lower divergence residuals than Binning, although its residual decreases slightly as k_n increases. This decrease does not indicate improved reconstruction accuracy, instead, as the measurement data become sparse, the PDE constraints and the network’s feature to generate smooth features exert a relatively stronger influence, resulting in smoother and more nearly divergence free fields despite simultaneous increase in velocity and vorticity error.

Figure 7 and Figure 8 present representative orthogonal slice views of the reconstructed velocity and pressure fields, together with the corresponding absolute-error fields, at $k_n = 4$ and $k_n = 64$, respectively. At $k_n = 4$, corresponding to a mean inter-particle spacing of approximately $9.5 l_\eta$, all three methods generate the dominant




Fig. 6 Normalized root-mean-square errors (NRMSEs) of (a) velocity, (b) pressure, (c) vorticity, and (d) divergence residual for the HIT fields reconstructed using Binning, FlowFit, and PINN at different particle down-selection factors k_n .

large-scale features of the velocity components. However, Binning exhibits stronger smoothing and higher deviations in the fine-scale structures. FlowFit preserves considerably more spatial detail, but the absolute error maps show local differences that are not apparent from the reconstructed fields alone. PINN produces the smallest and most spatially uniform error level for all three velocity components.

At $k_n = 64$, corresponding to a mean inter-particle spacing of approximately $24 l_\eta$, the fine-scale structures are substantially weakened for all reconstruction methods. Binning and FlowFit recover only the broad spatial variation of the flow, while displaying large localized velocity and pressure errors. PINN preserves a closer approximation of the dominant large-scale organization and produces lower errors than the other methods. Nevertheless, the PINN result also lacks a substantial fine-scale content in the ground truth. Its lower error should therefore be interpreted as improved reconstruction under severe data sparsity rather than complete recovery of the instantaneous small-scale flow field.

Figure 9 compares iso-surfaces of the Q-criterion at $Q = 3,500 \text{ s}^{-2}$, colored by the velocity magnitude. At $k_n = 1$ and $k_n = 4$, FlowFit and PINN preserve a large fraction of the elongated vortical structures observed in the ground truth, whereas Binning misses most of the small-scale structures. At $k_n = 16$, all methods exhibit a reduction in the number and continuity of structures, although the dominant structures remain




Fig. 7 Orthogonal domain-boundary plane slices of velocity and pressure, together with their absolute reconstruction errors, at the central time step of the HIT case for $k_n = 4$ ($l_p/l_\eta = 9.5$).




Fig. 8 Orthogonal domain-boundary plane slices of velocity and pressure, together with their absolute reconstruction errors, at the central time step of the HIT case for $k_n = 64$ ($l_p/l_\eta = 24$).




Fig. 9 Iso-surfaces of the Q-criterion at $Q = 3,500 \text{ s}^{-2}$, colored by the velocity magnitude, for the fields reconstructed using Binning, FlowFit, and PINN at different particle down-selection factors k_n .

for FlowFit and PINN. At $k_n = 64$, only a few structures in the Binning and FlowFit fields can be seen, while PINN retains a larger subset of the dominant vortical structures.

6.3 Synthetic TBL result

As in the HIT case, Figure 10 (a)–(d) presents the NRMSEs of velocity, pressure, and vorticity, together with the normalized divergence residual, for the synthetic TBL case. The velocity and pressure errors exhibit trends similar to those observed in the HIT case. For all three reconstruction methods, both velocity and pressure errors increase with the down-selection factor, as expected from the decreasing number of available particles. FlowFit consistently yields lower velocity and pressure errors than Binning across all down-selection factors and approaches the error level of PINN under densely seeded conditions. Nevertheless, PINN achieves the lowest errors over the entire




Fig. 10 Normalized root-mean-square errors (NRMSEs) of (a) velocity, (b) pressure, (c) vorticity, and (d) divergence residual for the synthetic TBL fields reconstructed using Binning, FlowFit, and PINN at different particle down-selection factors k_n .

range of down-selection factors. Its advantage becomes particularly pronounced under sparsely seeded conditions, demonstrating its robustness to particle-data sparsity.

For vorticity, FlowFit yields the lowest errors at several low-to-moderate down-selection factors, particularly at $k_n = 1, 2$, and 8 . At these conditions, FlowFit outperforms PINN in vorticity reconstruction despite exhibiting a larger velocity error. In contrast, the vorticity error of PINN varies less strongly with the down-selection factor, indicating more stable reconstruction performance as the particle density decreases. This behavior is also reflected in the Q-criterion iso-surfaces presented in Figure 13. Due to the divergence-free formulation of FlowFit, its divergence residual remains close to machine precision. The divergence residual of PINN is substantially lower than that of Binning but, as in the HIT case, decreases as the down-selection factor increases. This trend demonstrates that a smaller physics residual does not necessarily imply a more accurate reconstruction of the reference flow. In particular, increasingly smooth reconstructed fields may satisfy incompressibility while losing small-scale flow information. Unlike in the HIT case, the continuity residual of Binning varies non-monotonically with the down-selection factor. Because the present Binning formulation does not enforce incompressibility, increasing the number of particles per cell does not guarantee a more divergence-free reconstruction; the calculated divergence remains sensitive to the local particle distribution and the differentiation of the fitted velocity field.




Fig. 11 Orthogonal domain-boundary plane slices of velocity-fluctuation and pressure, together with their absolute reconstruction errors, at the central time step of the synthetic TBL case for $k_n = 4$ ($l_p/l_\nu = 9.5$).




Fig. 12 Orthogonal domain-boundary plane slices of velocity-fluctuation and pressure, together with their absolute reconstruction errors, at the central time step of the synthetic TBL case for $k_n = 64$ ($l_p/l_\nu = 24$).

Figure 11 and Figure 12 present orthogonal slices of the reconstructed velocity-fluctuation and pressure fields, together with the corresponding absolute-error fields, at $k_n = 4$ and $k_n = 64$, respectively. Unlike in the HIT case, velocity fluctuations instead of total velocities are shown to improve the visibility of differences between the reconstruction methods. The $k_n = 4$ case shown in Figure 11 corresponds to a normalized mean inter-particle spacing of approximately $l_p/l_\nu = 9.5$ and a seeding density of approximately 0.07 ppp under the imaging assumptions described above. This represents a practically relevant seeding condition for STB measurements. Even under this condition, Binning produces noticeably noisier velocity-fluctuation fields than the other two methods, particularly for the streamwise component near the wall. This behavior is likely related to the limited local particle support in a region characterized by strong velocity gradients and smaller flow scales than those in the outer region. FlowFit and PINN produce considerably smoother fields that more closely resemble the DNS reference. However, the absolute-error slices reveal that the PINN errors are generally lower than those of FlowFit, consistent with the quantitative results in Figure 10.

At $k_n = 64$, representing the most sparsely seeded synthetic case, PINN continues to preserve coherent structures in all three velocity-fluctuation components more successfully than the other methods. FlowFit develops pronounced near-wall noise, resembling the artifacts already observed for Binning at $k_n = 4$. Nevertheless, its reconstructed pressure field remains visually coherent, although the quantitative results in Figure 10 indicate a larger pressure error than that of PINN. The near-wall artifacts in the Binning reconstruction become more severe and are also evident in the spanwise velocity-fluctuation component. In addition, Binning exhibits substantial pressure deviations over the streamwise interval $0 - 0.02$ m. The corresponding absolute-error maps provide a more detailed comparison of these reconstruction differences.

Figure 13 presents iso-surfaces of the Q-criterion at $Q = 2.5 \times 10^6 \text{ s}^{-2}$, colored by velocity magnitude. Even in the most densely seeded case, Binning primarily recovers the large-scale flow structures while many smaller features are absent. With increasing down-selection, the reconstructed structures become increasingly fragmented, and eventually even some of the large-scale motions are no longer recovered. FlowFit reconstructs the vortical structures well under densely seeded conditions, particularly at $k_n = 1$ and 4, but develops increasingly noticeable near-wall artifacts as the particle density decreases. These artifacts may be associated, at least in part, with the virtual particles introduced into the fitting dataset to impose the no-slip condition, although a dedicated sensitivity study would be required to establish this causality. PINN, despite imposing the no-slip boundary condition as a soft constraint, preserves coherent vortical structures over a considerably wider range of particle densities.

Finally, Figure 14 presents wall-normal profiles of the reconstructed velocity statistics in viscous units at $k_n = 4$, providing an assessment of the reconstruction performance over the complete time series. Because only 51 timesteps are available, the resulting statistics are not assumed to be fully converged. The reconstructed profiles are therefore compared directly with DNS statistics calculated over the same temporal window rather than being evaluated primarily against the canonical viscous-sublayer relation and logarithmic law. Despite the insertion of virtual particles to impose the




Fig. 13 Iso-surfaces of the Q-criterion at $Q = 2.5 \times 10^6 \text{ s}^{-2}$, colored by the velocity magnitude, for the fields reconstructed using Binning, FlowFit, and PINN at different particle down-selection factors k_n .

no-slip condition, the Binning profile in Figure 14 (a) does not fully approach the expected zero value at the wall. In contrast, both PINN and FlowFit reproduce the no-slip condition accurately. Binning also exhibits the largest deviations from the DNS reference across the remaining velocity-statistics profiles. FlowFit accurately reconstructs the wall-normal component shown in Figure 14 (b), but larger deviations are observed in the streamwise and spanwise components near the wall. The PINN profiles nearly overlap with the DNS reference over most of the investigated y^+ range and across all velocity components. This agreement demonstrates that PINN provides a statistically faithful reconstruction over the available temporal sequence and suggests an advantage of its global spatio-temporal formulation.

6.4 Experimental TBL result

Figure 15 presents the validation NRMSEs of velocity and acceleration, together with the normalized divergence residual, for the experimental TBL case. Unlike the two synthetic cases, the experimental case does not provide dense ground-truth fields.




Fig. 14 Wall-normal profiles of the Reynolds normal stresses $\langle u'u' \rangle^+$, $\langle v'v' \rangle^+$, and $\langle w'w' \rangle^+$ at $k_n = 4$, compared with the DNS reference. The Reynolds normal stresses are computed from only 50 time steps, and full statistical convergence is therefore not guaranteed.

Therefore, the scattered particle positions, velocities, and accelerations obtained from TrackFit were divided into training and validation sets, as described in Section 5.4, and the reconstruction errors were evaluated at the validation points. Although the TrackFit estimates are not equivalent to DNS ground truth and retain measurement and fitting uncertainties, the two synthetic test cases showed that the particle quantities obtained from TrackFit were generally more accurate than the corresponding reconstructed Eulerian fields. They therefore provide a practical reference for evaluating the experimental reconstructions (Schneiders and Sciacchitano 2017). Because reference pressure and vorticity fields are unavailable, acceleration is used as an additional validation quantity. The acceleration comparison provides an indirect assessment of the pressure gradients and the temporal and spatial derivatives of the reconstructed velocity field, although it should not be interpreted as a direct measure of the complete velocity-gradient error.

The velocity NRMSE shown in Figure 15 (a) follows a trend similar to those observed in the two synthetic test cases. PINN and FlowFit outperform Binning across




Fig. 15 Normalized root-mean-square errors (NRMSEs) of (a) velocity, (b) acceleration, and (c) divergence residual for the experimental TBL fields reconstructed using Binning, FlowFit, and PINN at different particle down-selection factors k_n .

all down-selection factors. Under the most densely seeded condition, PINN and FlowFit exhibit comparable velocity errors, whereas PINN yields lower errors under most of the remaining seeding conditions. Figure 15 (b) shows that PINN also achieves the lowest validation error for acceleration. This result is noteworthy because Binning reconstructs acceleration by directly interpolating the TrackFit acceleration data, while FlowFit incorporates the TrackFit acceleration estimates into its fitting procedure and evaluates the acceleration from the reconstructed pressure gradient and viscous term. In contrast, PINN does not use measured acceleration as a training target; instead, acceleration is inferred from its reconstructed spatio-temporal flow field and the governing equations. Due to its divergence-free formulation, the continuity residual of FlowFit again remains close to machine precision. PINN exhibits a larger residual because incompressibility is imposed as a soft constraint rather than being satisfied exactly by construction.

Figure 16 and Figure 17 present orthogonal slices of the reconstructed velocity-fluctuation and pressure fields at $k_n = 4$ and $k_n = 32$, respectively. Because dense




Fig. 16 Orthogonal domain-boundary plane slices of velocity-fluctuation and pressure at the central time step of the experimental TBL case for $k_n = 4$ ($l_p/l_\nu = 14.8$).

reference fields are unavailable for the experimental case, these figures provide a qualitative comparison of the reconstructed flow structures rather than absolute-error fields. The overall behavior is similar to that observed in the synthetic TBL case. At $k_n = 4$, as shown in Figure 16, Binning produces noticeable noise in the velocity-fluctuation fields, particularly in the streamwise component near the wall. FlowFit also exhibits near-wall artifacts, although they are less severe than those observed for Binning. In contrast, PINN produces smoother and more spatially coherent velocity-fluctuation fields throughout the measurement domain.

At $k_n = 32$, the near-wall artifacts become substantially more pronounced in both the Binning and FlowFit reconstructions. PINN, however, continues to preserve coherent flow structures across all three velocity-fluctuation components despite the low particle density, although some degradation is visible compared with the more densely seeded case. The pressure field reconstructed by Binning changes relatively




Fig. 17 Orthogonal domain-boundary plane slices of velocity-fluctuation and pressure at the central time step of the experimental TBL case for $k_n = 32$ ($l_p/l_\nu = 29.4$).

little between the two seeding conditions, and PINN similarly produces a qualitatively consistent pressure distribution. In contrast, the FlowFit pressure field changes substantially as the particle density decreases.

Finally, Figure 18 presents iso-surfaces of the Q -criterion at $Q = 1.25 \times 10^6 \text{ s}^{-2}$, colored by velocity magnitude. Under the imaging assumptions used to estimate the particle-per-pixel value, the seeding density of the experimental case at $k_n = 1$ is approximately equivalent to that of the synthetic TBL case at $k_n = 8$. The qualitative response of the reconstruction methods to decreasing particle density is therefore consistent with that observed in the synthetic test case. Even under the most densely seeded experimental condition, Binning produces relatively fragmented Q -criterion structures, and most of these structures disappear as the particle density decreases. FlowFit preserves more coherent structures than Binning but exhibits increasing degradation near the wall, consistent with the velocity slices shown in Figure 16 and Figure 17. PINN produces the most coherent and least seeding-sensitive




Fig. 18 Iso-surfaces of the Q-criterion at $Q = 1.25 \times 10^6 \text{ s}^{-2}$, colored by the velocity magnitude, for the fields reconstructed using Binning, FlowFit, and PINN at different particle down-selection factors k_n .

Q-criterion structures across the investigated conditions. Nevertheless, a visible degradation appears at $k_n = 32$, for which the particle density is only one-quarter of that in the synthetic TBL case at $k_n = 64$. Because no reference Q-criterion field is available for the experiment, these results should be interpreted as evidence of structural coherence and robustness to particle sparsity rather than as a direct assessment of absolute reconstruction accuracy.

6.5 Computation cost

Considering that accurate statistical estimates from LPT measurements may require processing more than one hundred-thousand-time steps, the computational efficiency of DA methods is an important practical consideration. In this section, we compare the wall-clock runtime of each method for the investigated test cases using the hardware employed in this study. Table 2 summarizes the total runtime and average runtime per time step for each method. For all test cases, the PINN computations were performed on the Computer for Advanced Research in Aerospace (CARA) cluster using an NVIDIA A100 GPU, whereas FlowFit3 and Binning were executed on an AMD Ryzen Threadripper 1950X 16-core CPU. Under these hardware configurations, PINN

Table 2 Computational runtime of the DA methods for the HIT and TBL test cases.

Test case	Method	Total runtime (min)	Average runtime per time step (s)	Hardware
HIT	PINN	282	331	NVIDIA A100 GPU
	FlowFit3	62	73	AMD Ryzen Threadripper 1950X (16-core CPU)
	Binning	18	21	AMD Ryzen Threadripper 1950X (16-core CPU)
TBL	PINN	437	513	NVIDIA A100 GPU
	FlowFit3	19	22	AMD Ryzen Threadripper 1950X (16-core CPU)
	Binning	12	14	AMD Ryzen Threadripper 1950X (16-core CPU)

required the longest runtime for all test cases, followed by FlowFit3 and Binning. PINN required substantially longer runtimes than both FlowFit3 and Binning for all test cases.

Interestingly, FlowFit3 and Binning required longer runtimes for the HIT case than for the TBL case, whereas the opposite trend was observed for PINN. We hypothesize that this difference is partly related to the Fourier feature mapping used in the PINN. Because HIT is statistically isotropic, its spectral characteristics are similar across spatial directions, allowing a common Fourier-feature bandwidth to effectively represent the dominant spatial scales. In contrast, TBL is strongly anisotropic and contains different characteristic length scales in the streamwise, wall-normal, and spanwise directions. Consequently, the selection of appropriate Fourier-feature bandwidths becomes more challenging for the TBL case, potentially resulting in slower convergence.

7 Conclusions

This study systematically evaluated PINNs for the reconstruction of Eulerian velocity and pressure fields from three-dimensional LPT data over a wide range of particle concentrations. The PINN reconstruction was compared with polynomial Binning and FlowFit using identical TrackFit-processed particle data. Two synthetic datasets, homogeneous isotropic turbulence and a turbulent boundary layer, provided complete DNS reference fields, while an experimental turbulent-boundary-layer dataset was used to assess the applicability of the methods under actual measurement conditions. The

open web portal [2nd LPT and DA challenges](#) provides shared benchmark datasets for extending such comparisons to a broader range of flow conditions (Leclaire et al. 2026).

For the synthetic HIT case, PINN yielded the lowest velocity, pressure, and vorticity errors over the investigated range of particle down-selection factors. All reconstruction methods progressively lost small scale information as the particle concentration decreased, but PINN retained the dominant flow structures and exhibited a slower deterioration of the reconstructed fields. A similar tendency was observed for the synthetic TBL cases. PINN produced the lowest velocity and pressure errors throughout the investigated range, while FlowFit achieved lower vorticity errors at several densely seeded conditions. As the particle concentration decreased, however, the vorticity error of PINN varied substantially less than that of the conventional reconstruction methods, and coherent near-wall flow structures were retained consistent in time and space over a wider range of seeding densities. The wall-normal velocity statistics e.g. of Reynolds stresses further showed close agreement between the PINN reconstruction and the DNS reference over the given temporal sequence.

The comparison also demonstrates that satisfaction of an individual physical constraint should not be interpreted independently as a measure of reconstruction accuracy. FlowFit satisfies incompressibility to nearly machine precision through its hard divergence-free formulation, whereas continuity is imposed as a soft constraint in the present PINN. Nevertheless, the lower velocity, pressure, and gradient-sensitive errors obtained with PINN under sparse conditions show that a small divergence residual alone does not guarantee recovery of the underlying flow information.

The experimental turbulent-boundary-layer results support the trends observed in the synthetic cases. Both PINN and FlowFit achieved lower validation velocity errors than Binning, while PINN provided the lowest errors under most of the reduced-seeding conditions. Moreover, PINN yielded the lowest validation error in particle acceleration even though acceleration measurements were not used as a training target. Together with the comparatively coherent velocity and Q-criterion fields observed at low particle concentrations, this result indicates that the reconstructed PINN field retains physically meaningful temporal and spatial derivative information. Because no dense reference velocity, pressure, or vorticity field is available for the experiment, however, the experimental results should be interpreted as evidence of generalization and robustness to particle sparsity rather than as a direct validation of the complete reconstructed flow field.

Overall, the results demonstrate that the present PINN framework provides a robust approach for LPT-based flow reconstruction, with its advantages becoming increasingly pronounced as the available particle information decreases. Nevertheless, several limitations remain. The reconstruction relies on particle trajectories preprocessed by TrackFit, which may be incomplete, retain trajectory-fitting uncertainties, and are not explicitly constrained to satisfy the governing physics. In addition, the present formulation is restricted to incompressible flows, and pressure is the only unmeasured physical quantity inferred beyond the measured velocity field. Future work will therefore focus on jointly reconstructing particle trajectories and Eulerian flow fields under governing-equation constraints, extending the framework

to infer additional hidden physical quantities such as temperature, and developing more computationally efficient training strategies for large-scale spatio-temporal reconstruction.

Acknowledgements. We are grateful for DLR/DAAD Research Fellowships. We would like to acknowledge Ke Zhou and Sam Grauer for providing the particle advection data for the synthetic HIT and TBL test cases as used in [Zhou et al. \(2024\)](#) and the DLR project ADaMant for providing funding for the experimental TBL test case ([Schröder et al. 2024](#)).

References

- Agüera N, Caferio G, Astarita T, et al (2016) Ensemble 3d ptv for high resolution turbulent statistics. *Measurement Science and Technology* 27(12):124011
- Ansar T, Ashraf WM (2025) Comparison of kolmogorov–arnold networks and multi-layer perceptron for modelling and optimisation analysis of energy systems. *Energy and AI* 20:100473
- Buchwald T, Schanz D, Godbersen P, et al (2026) Accurate estimation of velocity and acceleration, their statistics and position prediction in Lagrangian Particle Tracking. PREPRINT (Version 1) available at Research Square <https://doi.org/10.21203/rs.3.rs-10342778/v1>
- Cai S, Gray C, Karniadakis GE (2024) Physics-informed neural networks enhanced particle tracking velocimetry: An example for turbulent jet flow. *IEEE Transactions on Instrumentation and Measurement* 73:1–9
- Cao Y, Fang Z, Wu Y, et al (2019) Towards understanding the spectral bias of deep learning. *arXiv preprint arXiv:191201198*
- Clark Di Leoni P, Agarwal K, Zaki TA, et al (2023) Reconstructing turbulent velocity and pressure fields from under-resolved noisy particle tracks using physics-informed neural networks. *Experiments in Fluids* 64(5):95
- Elsinga GE, Scarano F, Wieneke B, et al (2006) Tomographic particle image velocimetry. *Experiments in fluids* 41(6):933–947
- Gesemann S, Huhn F, Schanz D, et al (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
- Glorot X, Bengio Y (2010) Understanding the difficulty of training deep feedforward neural networks. In: *Proceedings of the thirteenth international conference on artificial intelligence and statistics, JMLR Workshop and Conference Proceedings*, pp 249–256

- Godbersen P, Gesemann S, Schanz D, et al (2024) Flowfit3: Efficient data assimilation of lpt measurements. In: 21st International Symposium on the Application of Laser and Imaging Techniques to Fluid Mechanics, 216, pp 1–14
- Jeon YJ, 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
- Kasagi N, Nishino K (1991) Probing turbulence with three-dimensional particle-tracking velocimetry. *Experimental thermal and fluid science* 4(5):601–612
- Krishnapriyan A, Gholami A, Zhe S, et al (2021) Characterizing possible failure modes in physics-informed neural networks. *Advances in neural information processing systems* 34:26548–26560
- Leclaire B, Cornic P, Champagnat F, et al (2026) Lpt and da challenges anytime: an open web portal for continuous development and benchmarking of algorithms. In: 22nd International Symposium on the Application of Laser and Imaging Techniques to Fluid Mechanics, Lisbon, Portugal
- LeCun Y, Bottou L, Bengio Y, et al (1998) Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86(11):2278–2324
- Nabian MA, Gladstone RJ, Meidani H (2021) Efficient training of physics-informed neural networks via importance sampling. *Computer-Aided Civil and Infrastructure Engineering* 36(8):962–977
- Raissi M, Perdikaris P, Karniadakis GE (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
- Salimans T, Kingma DP (2016) Weight normalization: A simple reparameterization to accelerate training of deep neural networks. *Advances in neural information processing systems* 29
- Scarano F (2013) 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(5):70
- Schanz D, Jahn T, Schröder A (2022) 3d particle position determination and correction at high particle densities. In: *Proceedings of the 20th International Symposium on the Application of Laser and Imaging Techniques to Fluid Mechanics*, 214, pp 1933–1949

- Schneiders JF, Scarano F (2016) Dense velocity reconstruction from tomographic ptv with material derivatives. *Experiments in fluids* 57(9):139
- Schneiders JF, Sciacchitano A (2017) Track benchmarking method for uncertainty quantification of particle tracking velocimetry interpolations. *Measurement Science and Technology* 28(6):065302
- Schneiders JF, Dwight RP, Scarano F (2014) Time-supersampling of 3d-piv measurements with vortex-in-cell simulation. *Experiments in Fluids* 55(3):1692
- Schröder A, Schanz D (2023) 3d lagrangian particle tracking in fluid mechanics. *Annual Review of Fluid Mechanics* 55(1):511–540
- Schröder A, Schanz D, Geisler R, et al (2024) Near-wall flow features in zpg-tbl at various reynolds numbers using dense 3d lagrangian particle tracking. In: 21st International Symposium on Application of Laser and Imaging Techniques to Fluid Mechanics, Lisbon, Portugal.
- Sciacchitano A, Leclaire B, Schröder A (2025) On the accuracy of data assimilation algorithms for dense flow field reconstructions. *Experiments in Fluids* 66(2):42
- Steinfurth B, Hassanein A, Doan N, et al (2024) Physics-informed neural networks for dense reconstruction of vortex rings from particle tracking velocimetry. *Physics of Fluids* 36(9)
- Toscano JD, Käufer T, Wang Z, et al (2024) Inferring turbulent velocity and temperature fields and their statistics from lagrangian velocity measurements using physics-informed kolmogorov-arnold networks. *arXiv preprint arXiv:240715727*
- Vyas N, Morwani D, Zhao R, et al (2025) Soap: Improving and stabilizing shampoo using adam for language modeling. In: *International Conference on Learning Representations*, pp 93423–93444
- Wang H, Liu Y, Wang S (2022a) Dense velocity reconstruction from particle image velocimetry/particle tracking velocimetry using a physics-informed neural network. *Physics of fluids* 34(1)
- Wang S, Teng Y, Perdikaris P (2021a) Understanding and mitigating gradient flow pathologies in physics-informed neural networks. *SIAM Journal on Scientific Computing* 43(5):A3055–A3081
- Wang S, Wang H, Perdikaris P (2021b) On the eigenvector bias of fourier feature networks: From regression to solving multi-scale pdes with physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering* 384:113938
- Wang S, Yu X, Perdikaris P (2022b) When and why pinns fail to train: A neural tangent kernel perspective. *Journal of Computational Physics* 449:110768

- Wang S, Li B, Perdikaris P, et al (2026) Gradient alignment in physics-informed neural networks: A second-order optimization perspective. *Advances in Neural Information Processing Systems* 38:168482–168532
- Wu C, Zhu M, Tan Q, et al (2023) A comprehensive study of non-adaptive and residual-based adaptive sampling for physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering* 403:115671
- Zhou K, Grauer SJ (2026) On the joint estimation of flow fields and particle properties from lagrangian data. *Journal of Fluid Mechanics* 1041:A5
- Zhou K, Grauer S, Schanz D, et al (2024) Benchmarking data assimilation algorithms for 3d lagrangian particle tracking. In: *21st International Symposium on Applications of Laser and Imaging Techniques to Fluid Mechanics 2024*, 229, pp 1–22

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [GraphicalAbstract.png](#)