

Bayesian Estimation of Localization Uncertainty in Joint Particle–Flow Reconstruction

R. Tang¹, N.Q. DiPatri^{2,3}, M.J. Yamakaitis⁴, C.M. Harwood^{2,3}, P.M. Bardet⁴, S.J. Grauer^{1,*}

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

2: IIHR—Hydroscience and Engineering, University of Iowa, United States

3: Dept. of Mechanical Engineering, University of Iowa, United States

4: Dept. of Mechanical and Aerospace Engineering, George Washington University, United States

*Corresponding author: sgrauer@psu.edu

Keywords: Lagrangian particle tracking, Flow reconstruction, Data assimilation, Uncertainty, Fluid–structure interactions.

ABSTRACT

Physics-informed data assimilation (DA) for Lagrangian particle tracking (LPT) reconstructs velocity and pressure fields from particle tracks by balancing the data with prior information from the governing equations. This balance depends on the assumed particle localization uncertainty, whose depth-wise component, σ_z , is often poorly characterized. Mis-specified uncertainty can cause the reconstruction to overfit to noise or smooth out genuine fluctuations, producing errors that propagate to the pressure field and downstream quantities of interest. We address this problem by reformulating neural-implicit particle advection (NIPA), a DA method for LPT, as a Bayesian inference problem. The resulting framework infers σ_z together with the latent particle positions and Eulerian flow fields. Measurement-based uncertainty estimates are incorporated as weak priors on σ_z , and the extension has nearly the same computational cost as the original formulation. The method is evaluated using a synthetic turbulent channel flow, an experimental zero-pressure-gradient turbulent boundary layer, and an experimental fluid–structure interaction case. In the synthetic case, the inferred uncertainty recovers the true uncertainty to within 3% across a sixteen-fold sweep of prior centers. In the experimental cases, the inferred values of σ_z and reconstructed fields are stable across the same sweep. These results show that depth-wise localization uncertainty can be inferred using DA, reducing the sensitivity of reconstructions to manually prescribed uncertainties.

1. Introduction

Turbulent flows of practical interest are often analyzed through dense, time-resolved Eulerian fields of velocity, pressure, and gradients of the same. Such fields support, among other applications, the identification of coherent structures (Schanz et al., 2018; Schröder et al., 2009), the inference of pressure and aerodynamic loads (Rival & van Oudheusden, 2017; Jux et al., 2020), and the extraction of dynamically important modes in fluid–structure interactions (FSI, Young et al., 2023; Kulkarni et al., 2024). Scale-resolving simulations can provide these fields, but their cost remains

prohibitive in many cases, and it can be difficult to accurately model unsteady boundary conditions, multiphase coupling, chemistry, and other effects (Slotnick et al., 2014). Experiments are therefore needed to observe the actual behavior of complex flows. Among available diagnostics, Lagrangian particle tracking (LPT, Schröder & Schanz, 2023) is a leading approach for volumetric measurement of fluid flow, providing non-intrusive, time-resolved, three-dimensional (3D) data in the form of particle tracks. These tracks are spatially sparse and often noisy, and converting them into continuous Eulerian fields entails a reconstruction problem. Modern algorithms often pose this problem as a data assimilation (DA) task, in which LPT measurements are incorporated into a numerical solution of the governing equations (He et al., 2025; Zaki & Wang, 2025). The accuracy of such reconstructions depends on the solver fidelity as well as the assumed uncertainty of the particle positions, but the latter is often poorly characterized in LPT experiments.

Individual particles in LPT are localized in two-dimensional (2D) images, often with sub-pixel accuracy, and several methods are available for estimating the associated uncertainty (Wernet & Pline, 1993; Westerweel, 2000; Bhattacharya & Vlachos, 2020). Volumetric LPT lifts these 2D measurements into 3D particle positions using multi-view reconstruction, e.g., triangulation (Maas et al., 1993; Wieneke, 2013), or single-camera depth-sensing methods such as those based on digital in-line holography (DIH, Toloui & Hong, 2015; Toloui et al., 2017; Kumar & Hong, 2025), plenoptic imaging (Hall et al., 2017; Moaven et al., 2024; Yamakaitis, Rezaiguia, et al., 2025), defocus-based localization (Barnkob et al., 2015; Fuchs et al., 2023), or synthetic-aperture methods (Bajpayee, 2014). In each case, the depth coordinate must be inferred via the camera configuration and image-formation model. It is therefore useful to distinguish the relatively well characterized in-plane uncertainty, originating from 2D particle localization, from the out-of-plane, or “depth-wise,” uncertainty, which arises at the 3D localization stage.

To better understand this distinction, consider the simplified triangulation geometry illustrated in Fig. 1. A particle P is detected in two images acquired at the same time by cameras A and B . Based on the camera positions, marked as x_A and x_B , and their calibrations, the 3D particle position x_P is inferred from the 2D projections $s_P^{(A)}$ and $s_P^{(B)}$. These positions in the image planes can be determined with high accuracy. The depth direction is approximately aligned with the mean viewing axis of a multi-camera system, as shown in Fig. 1, or with the optical axis of a single-camera system. The in-plane directions in the measurement volume are normal to the depth direction and approximately aligned with the image planes. In this work, we refer to the depth direction as z , with x and y being in-plane directions. The in-plane uncertainties are σ_x and σ_y (or simply σ_{xy} when $\sigma_x \approx \sigma_y$) and the depth-wise uncertainty is σ_z .

As can be appreciated from Fig. 1, a simplified geometric treatment yields

$$\sigma_z = \frac{\sigma_{xy}}{\sin(\theta/2)},$$

where θ is the angle between views A and B . This model shows the depth position is expected to be the least reliable component, i.e., when the viewing angle is limited, small errors in the




Figure 1. Particle localization uncertainty. Left: triangulation from two views. Right: geometric amplification of image-plane uncertainty along the depth direction.

image plane are amplified in the depth direction. In practice, σ_z can exceed σ_{xy} by several factors, ranging from about $2\times$ for conventional multi-camera setups to nearly $10\times$ for some single-camera methods (Zhou & Grauer, 2026). Moreover, localization uncertainty is not determined solely by the optical system. It also depends on calibration uncertainties, the seeding density, reconstruction and tracking parameters, filtering, and other aspects of the data-processing workflow (Bhattacharya & Vlachos, 2020; Gesemann, 2021; Janke & Michaelis, 2021). The combined effects of these steps on localization errors are difficult to characterize. Consequently, although σ_{xy} provides a useful reference scale, σ_z is greater than σ_{xy} , it is less reliably known, and it is more consequential for the flow field reconstruction, as we show in this work.

Uncertainty quantification (UQ) is central to particle-based velocimetry, and UQ for particle image velocimetry (PIV) provides the most mature example (Sciacchitano, 2019). The literature on PIV UQ provides useful context for LPT, although the uncertainty of interest is different: PIV UQ typically concerns the displacement uncertainty, whereas LPT requires uncertainty estimates of particle positions along a track. We represent this uncertainty by the positional covariance matrix Σ_x . A comprehensive treatment was proposed by Bhattacharya & Vlachos (2020), who developed a framework for propagating image-plane, reprojection, and calibration uncertainties through the 3D particle localization process. Their method yields an estimate of Σ_x for each particle position, which can be further propagated to estimate the uncertainty in velocity and pressure. Hans et al. (2024) later developed a “Bayesian volumetric reconstruction” method that treats the particle positions as latent variables. A forward imaging model maps a candidate particle field to the camera images, and variational inference is used to estimate the posterior distribution over the particle positions. In the present work, we adopt a related Bayesian viewpoint, including the estimation of latent particle positions, but we apply it at the flow reconstruction stage.

A complementary line of work treats uncertainty at the track level, i.e., after the particle positions have been reconstructed and linked in time. This stage is important because velocity and accelera-

tion are obtained by differentiating particle positions, which can lead to the amplification of small localization errors. Gesemann et al. (2016) addressed this problem by fitting B-splines to noisy track data, with penalty terms used to suppress high-frequency artifacts in the particle positions. TrackFit extends this idea by selecting filtering and interpolation parameters from the spectral content of the raw particle tracks (Gesemann, 2021). The method estimates the power spectral density of the position signals, separates the low-frequency particle motion from the high-frequency noise floor, and chooses a filter that approximates an ideal Wiener filter. Error propagation is then used to estimate the uncertainty of derived quantities, namely: particle positions, velocities, and accelerations over time. Uncertainties estimated by TrackFit are examined in this work. Janke & Michaelis (2021) proposed a related method based on adaptive filtering via local regression along the tracks. Their method estimates position, velocity, and acceleration uncertainties along each track and uses these estimates to automatically select the local filter length, reducing user inputs.

The aforementioned studies estimate uncertainty at distinct junctions in the LPT pipeline. The first set of methods estimates uncertainty after 3D localization; the second set does so once the particles have been tracked. The resulting estimates of Σ_x can be used for DA, but two questions arise. First, we ask how accurate are these uncertainty estimates? Rajendran et al. (2021) addressed this issue for PIV through a meta-uncertainty framework, which characterizes the sensitivity and reliability of competing uncertainty estimators. They found that the performance of individual methods can vary substantially across experiments, and they showed how these differences can be combined to estimate the uncertainty more robustly. The relative paucity of UQ methods for LPT, however, limits benefits of such an approach. A second question to ask is: what are the implications of Σ_x for downstream quantities of interest (QoI)? In PIV, Sciacchitano & Wieneke (2016) propagated uncertainty to derived quantities such as vorticity, velocity divergence, mean velocity, and Reynolds stresses, showing the strong dependence of these QoI on spatial error correlations and finite-sampling effects. Zhang et al. (2020) later showed how uncertainty estimates can be used to improve the accuracy of pressure estimates. These works demonstrate the material effects of uncertainty on velocity gradients, pressure fields, and other QoI in PIV processing.

Measurement uncertainty is similarly important in DA methods for LPT because Σ_x controls the weight assigned to the tracks relative to the prior information that regularizes the reconstruction. Depending on the formulation, this information may be imposed through control variables that determine the evolution of the flow field, such as parameterized initial and boundary conditions when a high-fidelity solver is used to enforce the governing equations (Zaki & Wang, 2025). In softly constrained formulations, the reconstruction may instead be regularized through penalties on governing equation residuals (He et al., 2025). Either way, Σ_x determines whether fluctuations along a measured track are interpreted as genuine particle motion or as noise. Underestimating the uncertainty can yield reconstructions that are overfitted to noise, injecting erroneous high-frequency fluctuations into the velocity field; overestimating the uncertainty weakens the data constraint and can smooth out real small-scale motions. Although these failure modes are well

known, UQ analyses for DA reconstructions in LPT remain underdeveloped.

The present work addresses the need for a DA framework in LPT that accounts for uncertainty in Σ_x . We do so by reformulating neural-implicit particle advection (NIPA, Zhou & Grauer, 2026) as a Bayesian estimation problem that simultaneously estimates the latent particle positions, flow fields, and depth-wise uncertainty, σ_z , from LPT data. The approach is evaluated on three datasets: a synthetic turbulent channel flow case from the Johns Hopkins Turbulence Database (JHTDB, Perlman et al., 2007), near-wall plenoptic LPT measurements of a zero-pressure-gradient (ZPG) turbulent boundary layer (TBL), and an FSI case involving a flexible strut in the wake of a cylindrical bluff body. Across these cases, the inferred σ_z is stable across a sixteen-fold sweep of the prior center and reduces the sensitivity of the reconstructed flow fields to manually prescribed uncertainty estimates. The remainder of the paper is organized as follows. Section 2 presents the Bayesian formulation of NIPA and relevant uncertainty diagnostics. Section 3 describes the test cases, Sec. 4 summarizes the implementation, Sec. 5 reports the results, and Sec. 6 concludes the manuscript.

2. Methodology

Below, we reformulate NIPA as a Bayesian inference problem. The original NIPA method is described by Zhou & Grauer (2026), and its extension to FSI is reported by Tang et al. (2025). Here, we briefly summarize the flow, particle, and structure models used in NIPA in Sec. 2.1. We then develop the Bayesian formulation used to estimate the model parameters in Sec. 2.2. Posterior diagnostics for estimating and assessing σ_z are introduced in Sec. 2.3.

2.1. Flow, structure, and particle models

Our framework combines models of the flow, particles, and any solid structures, denoted by F , P , and S , respectively. The model parameters are estimated such that their outputs match the LPT data and are consistent with the relevant governing equations and known boundary conditions.

To start, let $\mathcal{V}_t \subset \mathbb{R}^3$ denote the measurement volume at time t , where the subscript indicates that the spatial domain may vary in time, as in FSI problems. The corresponding domain is

$$\Omega = \bigcup_{t \in \mathcal{T}} \mathcal{V}_t \times \{t\} \subset \mathbb{R}^4, \quad (1)$$

where $\mathcal{T}$ is the temporal domain of the measurement. The Eulerian flow field is represented by a coordinate-based neural network,

$$F : \Omega \rightarrow \mathbb{R}^4, \quad (\mathbf{x}, t) \mapsto (\mathbf{u}, p), \quad (2)$$

which maps a space–time input coordinate to the velocity $\mathbf{u}$ and pressure p . Additional fields, such as density or temperature, can be appended to the network outputs as needed. The trainable weights and biases of F are collected in θ_F .

The measured LPT data consist of particle positions $\mathbf{x}_i^{(k)}$ at measurement times $t_i^{(k)}$, where $k \in \{1, \dots, n_p\}$ indexes a track and $i \in \{1, \dots, n_k\}$ indexes the samples along that track. In NIPA, each track is represented by a “kinematics-constrained track” (KCT) model. For the k th track,

$$\mathbf{P}^{(k)} : \mathcal{T}^{(k)} \rightarrow \mathbb{R}^3, \quad t \mapsto \mathbf{v}^{(k)}, \quad (3)$$

where $\mathcal{T}^{(k)} \subseteq \mathcal{T}$ is the time span of the track. The trainable parameters $\theta_p^{(k)}$ define a continuous velocity history for the particle. Meanwhile, the *latent* (i.e., estimated) particle positions at the measurement times are represented using a hat symbol: $\hat{\mathbf{x}}_i^{(k)}$. The KCT model enforces kinematic consistency by requiring that

$$\hat{\mathbf{x}}_i^{(k)} - \hat{\mathbf{x}}_{i-1}^{(k)} = \int_{t_{i-1}^{(k)}}^{t_i^{(k)}} \mathbf{v}^{(k)} dt, \quad i = 2, \dots, n_k. \quad (4)$$

In other words, the estimated particle positions and velocity history are inter-related: changing $\theta_p^{(k)}$ changes the continuous trajectory between measured samples, and the latent positions $\hat{\mathbf{x}}_i^{(k)}$ anchor the trajectory at the measurement times. This model is illustrated in Fig. 2. The latent particle positions may be estimated from the measured positions subject to the localization covariance Σ_x . For inertial particles, additional parameters (particle size, density, etc.) can also be inferred from the LPT data (Zhou & Grauer, 2026), but the particles are treated as tracers in this work.




Figure 2. “KCT” model for particle trajectories. Curves show trajectories generated by different velocity parameters; color indicates particle speed. Markers denote prescribed particle positions, and the dotted curve indicates the ground truth. Equation (4) is a hard constraint, so all tracks pass through the prescribed positions and $\theta_p^{(k)}$ is unconstrained.

The wall geometry and velocity of solid bodies enter through a structure model, S . Let $\partial\mathcal{V}_t$ denote the domain boundary and let $\Gamma_t \subseteq \partial\mathcal{V}_t$ be the portion occupied by the structure. The role of S is to provide a time-continuous, differentiable description of the structure kinematics, from which Γ_t and the wall velocity can be evaluated during training. The TBL case has a fixed wall, which

is trivial to model. The FSI case involves a flexible strut, whose geometry is illustrated in Fig. 3. The strut is a slender body with an elliptical cross-section and a spanwise coordinate s . We parameterize its motion using a transverse deflection δ and a twist angle β about the spanwise axis. The cross-section has semi-axes A and B , so the strut surface can be sampled from A , B , and the instantaneous profiles of δ and β along s .




Figure 3. Strut geometry. Left: undeformed strut. Middle: deformed configuration with transverse bending δ . Right: top view of the elliptical cross-section, with semi-axes A and B and twist β .

The strut kinematics are modeled by a neural network,

$$S : \mathbb{R}^+ \times \mathcal{T} \rightarrow \mathbb{R} \times [0, 2\pi), \quad (s, t) \mapsto (\delta, \beta). \quad (5)$$

The outputs of S are used to sample points $\mathbf{x}_s \in \Gamma_t$ on the strut surface. Moreover, the wall velocity needed for the moving-wall boundary condition is obtained by automatic differentiation, i.e., $\mathbf{w} = d\mathbf{x}_s/dt$ at the sampled locations. The parameters of S could be inferred together with the particle positions and flow field, as in Tang et al. (2025). However, in this case, high-fidelity measurements of the strut motion are available, as described in Sec. 3.3. We therefore train S in advance to and keep it fixed during the Bayesian reconstruction.

2.2. Bayesian formulation

The reconstruction begins with the particle positions measured by LPT (the “data”),

$$\mathcal{D} = \left\{ \mathbf{x}_i^{(k)} \mid k = 1, \dots, n_p, i = 1, \dots, n_k \right\},$$

where, again, k indexes the particle track and i indexes the measurements along it. The flow, particle, and structure models introduced above define continuous representations of the velocity and pressure fields, the particle trajectories, and the (potentially transient) domain boundaries. The flow model is parameterized by the weights and biases in θ_F ; the particle models are parameterized

by the KCT coefficients $\theta_p^{(k)}$ and the latent particle positions $\hat{x}_i^{(k)}$. The localization uncertainty is represented by the positional covariance matrix

$$\Sigma_x = \text{diag}(\sigma_x^2, \sigma_y^2, \sigma_z^2),$$

where we have assumed uncorrelated localization errors. We collect the unknowns as

$$\mathcal{Z} = \left\{ \theta_F, \theta_p^{(k)}, \hat{x}_i^{(k)}, \sigma \mid k = 1, \dots, n_p, i = 1, \dots, n_k \right\}, \quad \sigma = (\sigma_x, \sigma_y, \sigma_z).$$

The goal is to infer $\mathcal{Z}$ such that the latent particle positions are consistent with the measured LPT data, while the reconstructed tracks and flow fields satisfy the governing equations and boundary conditions. We formulate this task as a maximum a posteriori (MAP) estimation problem. Section 2.2.1 introduces the Bayesian objective, Sec. 2.2.2 specifies the likelihood, Sec. 2.2.3 defines the priors, and Sec. 2.2.4 summarizes the resulting loss function.

2.2.1. Brief primer on Bayesian inference

In Bayesian inference, the unknown QoI and the data are modeled as random variables. The *prior distribution* encodes what is known about the QoI before the data are observed and the *likelihood* describes the probability of observing said data for a given value of the QoI. Bayes' equation combines these terms to obtain the *posterior distribution*, which represents one's updated knowledge after a measurement. As mentioned above, the QoI are collected in $\mathcal{Z}$, the measured particle positions reside in $\mathcal{D}$, and $f(\cdot)$ denotes a probability density function (PDF). The posterior density is

$$f(\mathcal{Z} \mid \mathcal{D}) = \frac{f(\mathcal{D} \mid \mathcal{Z}) f_{\text{pr}}(\mathcal{Z})}{f(\mathcal{D})} \propto f(\mathcal{D} \mid \mathcal{Z}) f_{\text{pr}}(\mathcal{Z}). \quad (6)$$

Here, $f(\mathcal{D} \mid \mathcal{Z})$ is the likelihood, $f_{\text{pr}}(\mathcal{Z})$ is the prior density, and $f(\mathcal{D})$ is the evidence. Since the evidence is independent of $\mathcal{Z}$, it does not affect the reconstruction. The most probable set of QoI, called the MAP estimate, is obtained by maximizing the posterior density, or equivalently, by minimizing the negative log posterior,

$$\begin{aligned} \mathcal{Z}_{\text{MAP}} &= \arg \max_{\mathcal{Z}} [f(\mathcal{Z} \mid \mathcal{D})] \\ &= \arg \min_{\mathcal{Z}} \{-\log[f(\mathcal{Z} \mid \mathcal{D})]\} = \arg \min_{\mathcal{Z}} \{-\log[f(\mathcal{D} \mid \mathcal{Z})] - \log[f_{\text{pr}}(\mathcal{Z})]\}. \end{aligned} \quad (7)$$

The first term on the right-most side enforces consistency with the LPT data under the assumed localization uncertainty. The second term incorporates prior information through the governing equations, boundary conditions, and assumptions about the localization uncertainty.

2.2.2. Likelihood distribution

The measured particle positions are modeled as noisy observations of the true (latent) positions. We assume additive, zero-mean Gaussian localization errors with covariance Σ_x . For the k th particle at the i th measurement time, the measured position is modeled as

$$\mathbf{x}_i^{(k)} = \hat{\mathbf{x}}_i^{(k)} + \boldsymbol{\epsilon}_i^{(k)}, \quad \boldsymbol{\epsilon}_i^{(k)} \sim \mathcal{N}(\mathbf{0}, \Sigma_x),$$

where $\hat{\mathbf{x}}_i^{(k)}$ is the estimated position and $\boldsymbol{\epsilon}_i^{(k)}$ is the localization error. Assuming that such errors are independent, the likelihood is the product of likelihoods for all the particle positions,

$$f(\mathcal{D} \mid \mathcal{Z}) = \prod_{k=1}^{n_p} \prod_{i=1}^{n_k} \frac{1}{(2\pi)^{3/2} \det(\Sigma_x)^{1/2}} \exp \left[-\frac{1}{2} \left(\mathbf{x}_i^{(k)} - \hat{\mathbf{x}}_i^{(k)} \right)^\top \Sigma_x^{-1} \left(\mathbf{x}_i^{(k)} - \hat{\mathbf{x}}_i^{(k)} \right) \right]. \quad (8)$$

Note that the localization uncertainty is inferred, i.e., Σ_x depends on $\boldsymbol{\sigma} \in \mathcal{Z}$. As a result, the Gaussian normalization in Eq. (8) depends on $\mathcal{Z}$, and the determinant terms must be retained in our MAP objective.

2.2.3. Prior distributions

The prior density has two parts. The first is a *physics prior*, which penalizes residuals of the governing equations for the fluid and particle phases, as well as residuals of any prescribed boundary conditions. The second is a prior on $\boldsymbol{\sigma}$. We assume these sources of information are independent, such that

$$f_{\text{pr}}(\mathcal{Z}) = f_{\text{pr}}^{\text{phys}}(\mathcal{Z}) f_{\text{pr}}^{\text{unc}}(\boldsymbol{\sigma}). \quad (9)$$

Physics prior. The physics prior reflects the expectation that the reconstructed flow field, particle trajectories, and moving boundaries should be consistent with the governing equations and known boundary conditions. These constraints cannot be perfectly satisfied by the neural network and track models, owing to their finite representational capacity and to imperfections in the data and optimization. Thus, when these models are used to evaluate the governing equations, residuals are inevitable. We treat these residuals as centered random variables, which leads to a MAP objective with the familiar quadratic penalty terms used in PINNs.

For the fluids considered here, we use the incompressible Navier–Stokes equations in Ω . The local residual vector $\mathbf{e}_F \in \mathbb{R}^4$ contains a scalar continuity residual and a vector momentum residual,

$$e_{\text{cnt}} = \nabla \cdot \mathbf{u}, \quad (10a)$$

$$\mathbf{e}_{\text{mom}} = \frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} + \nabla p - \nu \nabla^2 \mathbf{u}, \quad (10b)$$

where $\mathbf{u}$ and p are the outputs of F for some input coordinate $(\mathbf{x}, t)$, ν is the kinematic viscosity, ∇ is the spatial gradient, and ∇^2 is the spatial Laplacian. The local fluid residual vector is thus

$$\mathbf{e}_F = (e_{\text{cnt}}, e_{\text{mom}}), \quad (\mathbf{x}, t) \in \Omega. \quad (11)$$

Spatial and temporal derivatives in Eq. (10) are evaluated by automatic differentiation of the flow model F .

For tracer particles, the particle prior reduces to a kinematic consistency condition between the particle and the surrounding fluid parcel. For the k th particle, the corresponding residual is

$$\mathbf{e}_P^{(k)} = \mathbf{u}[\hat{\mathbf{x}}^{(k)}(t), t] - \mathbf{v}^{(k)}(t), \quad t \in \mathcal{T}^{(k)}. \quad (12)$$

Here, $\hat{\mathbf{x}}^{(k)}$ and $\mathbf{v}^{(k)}$ are obtained from the k th KCT model $P^{(k)}$ and $\mathbf{u}$ is obtained by evaluating F at the k th particle's position at time t , as computed via $P^{(k)}$. For inertial particles, Eq. (12) must be replaced by a dynamical relation, e.g., from the Maxey–Riley equation in Zhou & Grauer (2026).

Boundary conditions are enforced on the relevant portion of the domain boundary, $\Gamma_t \subseteq \partial\mathcal{V}_t$. The corresponding space–time boundary is

$$\partial\Omega_S = \bigcup_{t \in \mathcal{T}} \Gamma_t \times \{t\} \subset \mathbb{R}^3 \times \mathbb{R}.$$

Let $\mathbf{w}$ denote the wall velocity on Γ_t , as represented by the surface model S . The no-slip or moving-wall residual on the interface is

$$\mathbf{e}_S = \mathbf{u} - \mathbf{w}, \quad (\mathbf{x}, t) \in \partial\Omega_S. \quad (13)$$

For a stationary wall, $\mathbf{w} = \mathbf{0}$. For FSI problems, $\mathbf{w}$ is determined by S .

The residual fields are summarized by normalized residual energies,

$$E_F^2 = \frac{1}{|\Omega|} \int_{\Omega} \|\mathbf{e}_F(\mathbf{x}, t)\|_2^2 \, d\mathbf{x} \, dt, \quad (14a)$$

$$E_P^2 = \frac{1}{n_p} \sum_{k=1}^{n_p} \frac{1}{|\mathcal{T}^{(k)}|} \int_{\mathcal{T}^{(k)}} \|\mathbf{e}_P^{(k)}(t)\|_2^2 \, dt, \quad (14b)$$

$$E_S^2 = \frac{1}{|\partial\Omega_S|} \int_{\partial\Omega_S} \|\mathbf{e}_S(\mathbf{x}, t)\|_2^2 \, d\mathbf{x} \, dt, \quad (14c)$$

which quantify the degree to which the models are inconsistent with the governing equations and known boundary conditions. These integrals are approximated by Monte Carlo sampling; sampling $\partial\Omega_S$ follows the procedure described by Tang et al. (2025). We convert these residual measures into a prior density by assuming that small residual energies are more probable than large ones, with the acceptable level set by the scales σ_F , σ_P , and σ_S . A Gaussian model gives

$$f_{\text{pr}}^{\text{phys}}(\mathcal{Z}) = \frac{1}{(2\pi)^{3/2} \sigma_F \sigma_P \sigma_S} \exp\left(-\frac{E_F^2}{2\sigma_F^2} - \frac{E_P^2}{2\sigma_P^2} - \frac{E_S^2}{2\sigma_S^2}\right). \quad (15)$$

Here, σ_F , σ_P , and σ_S determine the relative influence of the fluid, particle, and boundary condition priors. Smaller values impose the corresponding constraint more strongly; larger values allow for larger residuals. The residual energies E_F , E_P , and E_S depend on the model outputs and therefore on the unknown parameters in $\mathcal{Z}$.

Localization uncertainty prior. Standard deviations of the localization errors are positive scale parameters. We assume that reference estimates are available for each component of σ , denoted $\tilde{\sigma}_\alpha$ for $\alpha \in \{x, y, z\}$. Since these quantities are scales, errors in $\tilde{\sigma}_\alpha$ are naturally described by multiplicative factors. Hence, it is reasonable to model $\log(\sigma_\alpha)$ as normally distributed about $\log(\tilde{\sigma}_\alpha)$. We thus place independent log-normal priors on the components of σ , leading to

$$f_{\text{pr}}^{\text{unc}}(\sigma) = \prod_{\alpha \in \{x, y, z\}} \frac{1}{(2\pi)^{1/2} \sigma_\alpha \tau_\alpha} \exp\left\{-\frac{[\log(\sigma_\alpha) - \log(\tilde{\sigma}_\alpha)]^2}{2\tau_\alpha^2}\right\}. \quad (16)$$

The prior is centered at the reference value in log space, so $\tilde{\sigma}_\alpha$ is the median of the prior for σ_α . The width τ_α controls the spread about $\tilde{\sigma}_\alpha$. Since σ_x and σ_y are relatively well characterized, τ_x and τ_y (hereafter, τ_{xy}) are chosen to be small. By contrast, σ_z is more uncertain, so τ_z is chosen to be larger. The sensitivity of the inferred depth uncertainty to $\tilde{\sigma}_z$ is assessed in Sec. 5.

2.2.4. MAP objective loss

The joint density obtained by multiplying Eqs. (8), (15), and (16) is a very high-dimensional function. As a result, although the full posterior contains all information about the reconstruction and its uncertainty, in principle, it is analytically and computationally intractable. Nevertheless, if $f(\mathcal{Z} \mid \mathcal{D})$ is smooth, unimodal, and strongly peaked, we can characterize the solution using a representative point estimate, i.e., the MAP. We obtain $\mathcal{Z}_{\text{MAP}}$ by minimizing the negative log posterior, which yields the objective function used for training.

Before presenting the MAP objective, we note that it is convenient to optimize the localization uncertainties through unconstrained variables, i.e., without enforcing $\sigma_\alpha > 0$. Specifically, we introduce ξ_α for $\alpha \in \{x, y, z\}$, such that

$$\sigma_\alpha = \tilde{\sigma}_\alpha \exp(\xi_\alpha) \quad \text{and} \quad \xi_\alpha = \log\left(\frac{\sigma_\alpha}{\tilde{\sigma}_\alpha}\right).$$

This parameterization enforces $\sigma_\alpha > 0$, and $\xi_\alpha = 0$ corresponds to the reference value. Applying the change of variables from σ_α to ξ_α requires

$$f(\xi_\alpha) = f(\sigma_\alpha) \left| \frac{d\sigma_\alpha}{d\xi_\alpha} \right| = f(\sigma_\alpha) \sigma_\alpha.$$

Thus, the log-normal prior on σ_α becomes a normal prior on ξ_α :

$$f_{\text{pr}}^{\text{unc}}(\xi) = \prod_{\alpha \in \{x, y, z\}} \frac{1}{(2\pi)^{1/2} \tau_\alpha} \exp\left(-\frac{\xi_\alpha^2}{2\tau_\alpha^2}\right) = \prod_{\alpha \in \{x, y, z\}} \frac{1}{(2\pi)^{1/2} \tau_\alpha} \exp\left[-\frac{1}{2\tau_\alpha^2} \log\left(\frac{\sigma_\alpha}{\tilde{\sigma}_\alpha}\right)^2\right]. \quad (17)$$

The normalization in this PDF does not depend on ξ , simplifying evaluation of the uncertainty prior in the objective.

Using the above transformation, taking the negative logarithm of the joint density, and dropping additive constants that do not depend on $\mathcal{Z}$ yields the MAP objective,

$$\mathcal{J}_{\text{MAP}} = \mathcal{J}_{\text{data}} + \mathcal{J}_{\text{pr}}^{\text{phys}} + \mathcal{J}_{\text{pr}}^{\text{unc}}. \quad (18)$$

The individual components of this loss are

$$\mathcal{J}_{\text{data}} = \frac{1}{2n_p} \sum_{k=1}^{n_p} \frac{1}{n_k} \sum_{i=1}^{n_k} \left(\mathbf{x}_i^{(k)} - \hat{\mathbf{x}}_i^{(k)} \right)^\top \Sigma_{\mathbf{x}}^{-1} \left(\mathbf{x}_i^{(k)} - \hat{\mathbf{x}}_i^{(k)} \right) + \frac{1}{2} \log[\det(\Sigma_{\mathbf{x}})], \quad (19a)$$

$$\mathcal{J}_{\text{pr}}^{\text{phys}} = \frac{E_{\text{F}}^2}{2\sigma_{\text{F}}^2} + \frac{E_{\text{P}}^2}{2\sigma_{\text{P}}^2} + \frac{E_{\text{S}}^2}{2\sigma_{\text{S}}^2}, \quad (19b)$$

$$\mathcal{J}_{\text{pr}}^{\text{unc}} = \sum_{\alpha \in \{x,y,z\}} \frac{1}{2\tau_{\alpha}^2} \log \left(\frac{\sigma_{\alpha}}{\tilde{\sigma}_{\alpha}} \right)^2. \quad (19c)$$

Here, $\Sigma_{\mathbf{x}}$ and σ_{α} are evaluated using the current value of ξ . The primary distinction between the present formulation and those of Zhou & Grauer (2026) and Tang et al. (2025) is that the localization uncertainty is inferred. The log determinant term in $\mathcal{J}_{\text{data}}$ must be evaluated and $\mathcal{J}_{\text{pr}}^{\text{unc}}$ encodes reference information about the localization uncertainties. These terms are cheap to compute and they enable the inference of σ together with the particle trajectories and flow field. Indeed, without the log determinant term and uncertainty prior, the objective admits a trivial failure mode in which the elements of σ go to infinity, causing the data term to vanish, and the reconstruction is readily driven toward an arbitrary admissible solution such as $\mathbf{u} = \mathbf{0}$ and $p = 0$.

2.3. Posterior diagnostics for the localization uncertainty

The diagnostics introduced here allow us to assess whether the inferred uncertainties are consistent with the prior–data balance implied by Bayes equation. MAP estimates of the localization uncertainties are denoted by $\sigma_{\alpha}^{\text{MAP}}$. These values are regularized by the uncertainty prior and thus reflect a compromise between the data residual and the prior center, $\tilde{\sigma}_{\alpha}$. As a consequence, we do not generally expect $\sigma_{\alpha}^{\text{MAP}}$ to equal the corresponding residual scale.

We define the residual scale $\sigma_{\alpha}^{\text{res}}$ in terms of the mean squared error (MSE),

$$\sigma_{\alpha}^{\text{res}} = \sqrt{\text{MSE}_{\alpha}}, \quad \text{where} \quad \text{MSE}_{\alpha} = \frac{1}{n_p} \sum_{k=1}^{n_p} \frac{1}{n_k} \sum_{i=1}^{n_k} \left(x_{\alpha,i}^{(k)} - \hat{x}_{\alpha,i}^{(k)} \right)^2. \quad (20)$$

Here, $x_{\alpha,i}^{(k)}$ is the α -component of the measured position for the k th particle at the i th time index and $\hat{x}_{\alpha,i}^{(k)}$ is our estimate of the true coordinate. Given a perfect reconstruction and unbiased particle positions, $\sigma_{\alpha}^{\text{res}}$ should approach the true value, $\sigma_{\alpha}^{\text{GT}}$. Therefore, $\sigma_{\alpha}^{\text{res}}$ serves as an estimator of $\sigma_{\alpha}^{\text{GT}}$ that does not depend upon the prior.

Next, we define a scale-balance factor,

$$\kappa_\alpha = \left(\frac{\sigma_\alpha^{\text{res}}}{\sigma_\alpha^{\text{MAP}}} \right)^2. \quad (21)$$

This factor compares the residual scale implied by the data with the corresponding MAP estimate. It is generally not equal to unity because $\sigma_\alpha^{\text{MAP}}$ is regularized by the uncertainty prior. The expected value of κ_α at the MAP can be derived by considering the terms of $\mathcal{J}_{\text{MAP}}$ that explicitly depend on σ_α . Up to additive constants independent of σ_α , these terms are

$$\mathcal{J}_\alpha = \frac{\text{MSE}_\alpha}{2\sigma_\alpha^2} + \log(\sigma_\alpha) + \frac{1}{2\tau_\alpha^2} \log\left(\frac{\sigma_\alpha}{\tilde{\sigma}_\alpha}\right)^2. \quad (22)$$

The first term comes from the weighted particle position residual, the second from the Gaussian normalization in the likelihood, and the third from the uncertainty prior. At the MAP estimate, the gradient of $\mathcal{J}_{\text{MAP}}$ with respect to σ_α should vanish. Therefore, we set $\partial \mathcal{J}_\alpha / \partial \sigma_\alpha = 0$ at $\sigma_\alpha^{\text{MAP}}$ to obtain the fixed-point relation

$$\kappa_\alpha^{\text{MAP}} = 1 + \frac{1}{\tau_\alpha^2} \log\left(\frac{\sigma_\alpha^{\text{MAP}}}{\tilde{\sigma}_\alpha}\right). \quad (23)$$

The logarithmic term measures the shift of $\sigma_\alpha^{\text{MAP}}$ from the prior center, $\tilde{\sigma}_\alpha$. This term vanishes when the MAP estimate coincides with the reference value, in which case the fixed-point condition gives $\kappa_\alpha^{\text{MAP}} = 1$. If $\tilde{\sigma}_\alpha$ is mis-specified, however, then $\sigma_\alpha^{\text{MAP}}$ shifts away from $\tilde{\sigma}_\alpha$ and $\kappa_\alpha^{\text{MAP}} \neq 1$.

Agreement between the inferred value of κ_α and that predicted by Eq. (23) indicates convergence to $\mathcal{Z}_{\text{MAP}}$. However, this diagnostic does not, by itself, establish the accuracy of σ_α . For the synthetic results in Sec. 5.2, the accuracy of σ_α can be checked against the ground truth, but for the experimental cases, no such ground truth is available. Instead, we check whether κ_α agrees with Eq. (23), whether $\sigma_\alpha^{\text{res}}$ is robust to changes in $\tilde{\sigma}_\alpha$, and whether the inferred uncertainties are consistent with a priori estimates based on the imaging geometry.

3. Test cases

We evaluate the proposed framework on three datasets. The first case is a synthetic LPT dataset generated using direct numerical simulation (DNS) flow fields from the JHTDB. The second and third are experimental datasets measured in the Boundary Layer Wind Tunnel at the Naval Surface Warfare Center, Carderock Division, and in the Recirculating Water Channel facility at the University of Iowa, respectively. Figure 4 shows both experimental LPT setups, and the following subsections describe all three test cases.




Figure 4. Experimental setups for the LPT measurements. Left: ZPG-TBL experiment with side illumination and plenoptic imaging from below. Right: FSI experiment with a flexible strut downstream of a cylindrical bluff body and four-view imaging from below.

3.1. Turbulent channel flow simulation

A DNS of fully developed turbulent channel flow with a friction Reynolds number of 1000 is used for the synthetic case. Data for this flow are obtained from the JHTDB. The dimensionless domain spans $8\pi \times 3\pi \times 2$ and is discretized with $2048 \times 1536 \times 512$ voxels in the streamwise, spanwise, and wall-normal directions, with Chebyshev-type clustering toward the walls. The data are stored at a dimensionless time step of 0.0065. We extract two sub-volumes of identical streamwise and spanwise extent, with an aspect ratio matched to the FSI volume described in Sec. 3.3. The wall-normal thickness is fixed at $\Delta z = 0.20$, corresponding to $\Delta z^+ = 200$, so that each sub-volume measures roughly $1225 \times 833 \times 200$ in viscous units. Instantaneous w -velocity flow fields from these cases are shown in Fig. 5. Assuming air as the working fluid, with $\nu = 15 \text{ mm}^2/\text{s}$, this corresponds to a physical region of approximately $314 \times 225 \times 54 \text{ mm}^3$, where the viscous scales are $\ell_\nu = 0.25 \text{ mm}$ and $\tau_\nu = 4.3 \text{ ms}$.

The two sub-volumes differ only in their wall-normal placement. A “high-turbulence” sub-volume lies within the logarithmic layer, with $z^+ \in [52, 254]$, wherein the flow contains attached-eddy-type and hairpin-like vortical motions, streamwise elongated large-scale structures, and smaller motions influenced by the buffer layer below. A “low-turbulence” sub-volume is flush with the channel centerline, corresponding to $z^+ \in [800, 1000]$; flow in this domain is characterized by weaker, more isotropic fluctuations associated with larger outer-layer structures. These differences are readily apparent in the flow fields shown in Fig. 5. Owing to the non-uniform wall-normal grid, the high- and low-turbulence domains contain $101 \times 137 \times 67$ voxels and $101 \times 137 \times 34$ voxels, respectively. This pairing brackets the range of turbulence intensities anticipated in the FSI experiment and allows the framework to be evaluated across distinct dynamical regimes without the added complexity of a moving boundary.




Figure 5. Representative instantaneous out-of-plane w -velocity fields in the low- and high-turbulence sub-volumes for the channel flow case.

Synthetic LPT data are generated by uniformly seeding 12,000 massless tracer particles in each sub-volume and advecting them through the DNS velocity field with a second-order Runge–Kutta integration scheme. The seeding density is chosen to match that of the FSI measurements reported in Sec. 3.3, resulting in a mean inter-particle spacing of $29\ell_\nu$. Periodic boundary conditions are applied in the streamwise and spanwise directions; particles that exit through the wall-normal faces of the sub-volume are reinjected at uniformly sampled wall-normal positions to maintain a statistically stationary distribution. The simulation is advanced over 50 snapshots using a time step of $\Delta t = 1.4$ ms.

To mimic the localization errors in the FSI experiment, the true particle positions are corrupted with additive, independent, zero-mean Gaussian noise. Anisotropic localization errors are introduced by setting the standard deviations to $\sigma_{xy} = 34 \mu\text{m}$ and $\sigma_z = 340 \mu\text{m}$, corresponding to 0.2 and 2 pixels, respectively, given the MiniShaker pixel pitch of 0.17 mm. This yields a true out-of-plane to in-plane uncertainty ratio of $\sigma_z/\sigma_{xy} = 10$. The in-plane uncertainty is supplied to the framework as a known reference value, anchored by a tight log-normal prior, and σ_z is either fixed or treated as unknown and inferred jointly with the flow and particle states.

3.2. Turbulent boundary layer experiment

Measurements of a ZPG-TBL were performed in a blow-down, open-circuit boundary layer wind tunnel (Yamakaitis, Pigeon, et al., 2025). The tunnel has a 3.66 m long test section with a $0.46 \times 0.91 \text{ m}^2$ cross-section, and the ceiling can be adjusted for pressure gradient control. The boundary layer develops along a smooth acrylic floor with a free-stream velocity of $U_\infty = 10 \text{ m/s}$, with the ceiling angle set to produce a ZPG condition. Measurements were taken 3.35 m downstream of the test section entrance, providing a long development length. At that point, the flow had a friction Reynolds number of $Re_\tau \approx 1640$, with viscous scales of $\ell_\nu = 46 \mu\text{m}$ and $\tau_\nu = 0.135 \text{ ms}$ and a friction velocity of $u_\tau = 0.342 \text{ m/s}$.

Imaging was performed with a custom “Plenoptic 3.0” system (Yamakaitis, Rezaiguia, et al., 2025),

pictured on the left side of Fig. 4, that produces four non-overlapping views on a single high-speed camera sensor, mimicking a conventional four-camera LPT arrangement. The camera was mounted below the test section to view the near-wall region through a window, while the measurement volume was illuminated from the side by a dual-head laser system (Photonics Industries DMD527-50) producing 200 ns pulses. The common measurement volume was a cylindrical region of 6 mm diameter and 1 mm wall-normal height. Images were recorded at 9.95 kHz. Particle trajectories were extracted using the Shake-the-Box (STB) algorithm in DaVis 10.2.1. The two runs considered here contained approximately 16.2×10^6 and 18.6×10^6 recovered tracks, respectively, with mean track lengths of 17.6 and 16.1 frames and approximately 478–500 tracks per frame. Further details of the setup and processing workflow are reported by Yamakaitis, Pigeon, et al. (2025).

For the present reconstruction, 250 consecutive frames were selected from one of the recordings, corresponding to approximately 25 ms of measurement. The per-particle localization uncertainty was estimated assuming a sub-pixel detection accuracy of 0.1 px, converted to physical units through the pixel pitch and magnification of the imaging system. This gives an in-plane uncertainty of $\sigma_{xy} \approx 3 \mu\text{m}$. As illustrated in Fig. 1, two viewing rays intersecting at angle θ , each having a lateral uncertainty of σ_{xy} , yield a depth-wise uncertainty of $\sigma_z \approx \sigma_{xy} / \sin(\theta/2)$ along the bisector (Yamakaitis, Rezaiguia, et al., 2025). Applying this relation to the plenoptic system's aperture gives $\sigma_z \approx 14\text{--}15 \mu\text{m}$, approximately five times σ_{xy} . In viscous units, this corresponds to $\sigma_{xy} \approx 0.07\ell_\nu$ and $\sigma_z \approx 0.30\text{--}0.33\ell_\nu$. These values are used as the prior estimates $\tilde{\sigma}_\alpha$ in our Bayesian reconstructions.

3.3. Fluid–structure interaction experiment

The experimental FSI case, shown on the right side of Fig. 4, features the interaction between a cylinder wake and a flexible strut in a recirculating water channel. The channel has a length of 10 m and cross-section of $0.61 \times 0.61 \text{ m}^2$. A rigid circular cylinder of diameter $D = 60 \text{ mm}$ is mounted upstream of the flexible PVC strut, with the strut's leading edge located $6.26D$ downstream of the cylinder axis. The cylinder produces a vortex street that periodically forces the strut. The strut has an elliptical cross-section of $57.15 \times 20.32 \text{ mm}$ and is instrumented with a “kinematic shape sensing” system (KSS system, Di Napoli & Harwood, 2020) that records its motion. Broadband excitation is supplied by turbulent inflow, and a shaker motor attached to the strut supplies additional mechanical excitation. The condition analyzed here corresponds to a free-stream velocity of $U = 0.385 \text{ m/s}$, giving $Re_D = 17,751$ based on the cylinder diameter.

The flow was seeded with $60 \mu\text{m}$ polyamide particles and illuminated bilaterally by pulsed LEDs. Particles were imaged from beneath the channel at 128 Hz by a four-view imager (LaVision MiniShaker TR), with the optical axis aligned with the strut span. The reconstructed measurement volume is roughly $314 \times 225 \times 54 \text{ mm}^3$ in size and contains around 90,000 particles per frame, yielding an image density of 0.042 ppp. Particle trajectories were extracted using STB in DaVis 10.2.1, which gave 12,000–16,000 tracks per snapshot. The structural response was independently mea-

sured using the embedded KSS sensor array, from which bending and torsional displacements were reconstructed. The KSS strain signals were sampled at 2048 Hz, and the full-rate structural data were used to train the structure model S.

For the FSI case, 50 consecutive time steps were analyzed, corresponding to approximately 0.39 s of measurement. The volume self-calibration achieved a root-mean-square disparity of 0.10 pixels at a pixel pitch of 0.17 mm, giving an in-plane localization uncertainty of $\sigma_{xy} \approx 17 \mu\text{m}$. TrackFit provided an initial depth-wise uncertainty estimate of $85 \mu\text{m}$. However, the camera calibration yielded separation angles of approximately $8\text{--}12^\circ$, which correspond to depth-wise uncertainties of $162\text{--}244 \mu\text{m}$ using the simplified geometric relation in Fig. 1. We therefore used an assumed uncertainty ratio of $\sigma_z/\sigma_{xy} = 10$, giving the intermediate reference value $\sigma_z^* = 170 \mu\text{m}$ (i.e., between the TrackFit and geometric estimates) for the prior sweep described in Sec. 5.1.

4. Implementation of the method and error metrics

The framework is implemented in TensorFlow 2. Exact partial derivatives of the model outputs and of the loss with respect to the trainable parameters are obtained by automatic differentiation. Integrals over Ω , $\partial\Omega_S$, and $\mathcal{T}^{(k)}$ in the physics prior are approximated by Monte Carlo sampling. The velocity and pressure fields are represented by two independent sub-networks, $F_u : (\mathbf{x}, t) \mapsto \mathbf{u}$ and $F_p : (\mathbf{x}, t) \mapsto p$. Both use the PirateNet architecture of Wang et al. (2024) with two residual blocks, swish activations, and Xavier-normal weight initialization. To mitigate the spectral bias of gradient-based training, we apply a Fourier encoding with 256 random frequencies sampled at standard deviations of 2 for spatial features and 1 for temporal features (Wang et al., 2021). In all cases, the hidden layers are 300-neurons wide for the velocity network and 150 wide for the pressure network.

A tight prior is placed on the in-plane localization uncertainties, with $\tau_{xy} = 0.3$, reflecting the fact that these uncertainties are comparatively well characterized. A broader prior, with $\tau_z = 2$, is used for the depth-wise uncertainty. These priors are discussed in more detail in Sec. 5.1. The residual scales σ_F , σ_P , and σ_S act as fixed weights in the physics prior and are chosen through a coarse parameter sweep. For the TBL case, we use

$$\sigma_F = 10^{-3/2}, \quad \sigma_P = 10^{-5/2}, \quad \text{and} \quad \sigma_S = 10^{1/2}.$$

For the FSI case, we use

$$\sigma_F = 10^{-1}, \quad \sigma_P = 10^{-2}, \quad \text{and} \quad \sigma_S = 10.$$

There is no wall in the channel flow sub-volumes, so only the fluid and particle residual scales are needed, with

$$\sigma_F = 10^{-1} \quad \text{and} \quad \sigma_P = 10^{-2}.$$

All trainable parameters are updated using the Adam optimizer at a learning rate of 10^{-3} for 2×10^5 iterations, with one full run taking approximately six hours on an NVIDIA L40S GPU.

Reconstruction accuracy can be quantified in the synthetic case. Flow field and particle position errors are reported using the normalized root-mean-square error (NRMSE),

$$e_\varphi = \left(\frac{\langle \|\varphi - \varphi_{\text{GT}}\|_2^2 \rangle_\star}{\langle \|\varphi_{\text{GT}}\|_2^2 \rangle_\star} \right)^{1/2},$$

where $\langle \cdot \rangle_\star$ denotes an average over a suitable domain. For flow fields, this average is taken over Ω ,

$$\langle \varphi \rangle_\Omega = \frac{1}{|\Omega|} \int_\Omega \varphi(\mathbf{x}, t) \, d\mathbf{x} \, dt.$$

For particle positions, the average is taken over the reconstructed tracks,

$$\langle \varphi \rangle_{\text{p}} = \frac{1}{n_{\text{p}}} \sum_{k=1}^{n_{\text{p}}} \frac{1}{|\mathcal{T}^{(k)}|} \int_{\mathcal{T}^{(k)}} \varphi(t) \, dt.$$

In practice, these integrals are approximated by sums over DNS grid points and time steps for flow fields and over the measurement instants in $\mathcal{T}^{(k)}$ for particle tracks. Flow field errors are evaluated using the fluctuating components after the Reynolds decomposition, $\varphi = \bar{\varphi} + \varphi'$. Lastly, localization uncertainty errors are reported as

$$e_{\sigma, \alpha} = \frac{|\sigma_\alpha^{\text{res}} - \sigma_\alpha^{\text{GT}}|}{\sigma_\alpha^{\text{GT}}}.$$

We focus on the error in the depth-wise estimate, $e_{\sigma, z}$.

5. Reconstruction results

We now present and analyze reconstructions for all three test cases. Section 5.1 defines the testing protocol. The channel flow results are then used to validate the method in Sec. 5.2, followed by its application to the experimental ZPG-TBL and FSI cases in Secs. 5.3 and 5.4, respectively.

5.1. Testing protocol

For each case, we compare two reconstruction strategies. First, we use fixed localization uncertainties, corresponding to the limit $\tau \rightarrow 0$. Second, we allow the uncertainties to vary, with $\tau_{xy} = 0.3$ and $\tau_z = 2$, corresponding to a tight prior on the in-plane uncertainties and a loose prior on the depth-wise uncertainty. These prior PDFs are shown in Fig. 6. The tight prior places the central 90% of its density on $\sigma_{xy}/\tilde{\sigma}_{xy} \in [0.82, 1.22]$, so σ_{xy} remains close to the reference scale, $\tilde{\sigma}_{xy}$. For the




Figure 6. Prior density for the normalized localization uncertainty, $\sigma_\alpha/\tilde{\sigma}_\alpha$. Curves show the tight prior used for in-plane uncertainties ($\tau_{xy}a = 0.3$) and the broad prior used for depth-wise uncertainty ($\tau_z = 2$). Vertical dotted lines mark the 5% and 95% percentiles as well as $\sigma_\alpha = \tilde{\sigma}_\alpha$.

loose prior, the corresponding interval is $[0.26, 3.85]$, thus allowing σ_z to vary much more freely. We refer to the $\tau \rightarrow 0$ limit as “fixed- σ ” reconstruction and to the $\tau_{xy} = 0.3$ and $\tau_z = 2$ configuration as “variable- σ ” reconstruction.

The *reference* in-plane uncertainty, $\tilde{\sigma}_{xy}$, is fixed in all tests, with $\tilde{\sigma}_{xy} = \sigma_{xy}^*$, where σ_α^* represents the true value in the synthetic case and our best a priori estimate in the experimental cases. The depth-wise reference value is varied to control the ratio

$$r_z = \frac{\tilde{\sigma}_z}{\sigma_z^*}.$$

For both the fixed- σ and variable- σ reconstructions, we consider

$$r_z \in \{0.25, 0.5, 1, 2, 4\}.$$

Thus, the prior value of σ_z is varied over a sixteen-fold range, with either a fixed ($\tau_z \rightarrow 0$) or variable ($\tau_z = 2$) configuration. These tests address two related questions: (1) Can σ_z be accurately inferred by MAP estimation? (2) How do the velocity and pressure field reconstructions respond when $\tilde{\sigma}_z$ is mis-specified?

For the experimental cases, the initial sweeps were performed using preliminary estimates of σ_z^* that were computed using TrackFit or based on the imaging geometry. After these sweeps, we revised the reference estimates using the diagnostics introduced in Sec. 2.3. We then performed one more variable- σ reconstruction for each case, setting $\tilde{\sigma}_z$ to the revised reference value whilst keeping $\tau_{xy} = 0.3$ and $\tau_z = 2$. These final reconstructions allow us to reassess the diagnostic criteria using updated uncertainty estimates.

5.2. Validation on the synthetic channel flow dataset

Results for the channel flow case are divided into two parts. First, the accuracy of Bayesian inference of σ_z is assessed in Sec. 5.2.1. Second, we examine how the assumed depth-wise uncertainty, $\tilde{\sigma}_z$, affects the reconstructed particle positions and flow fields in Sec. 5.2.2.

5.2.1. Estimation of out-of-plane localization uncertainty

We begin by focusing on the accuracy of σ_z inference in the variable- σ reconstructions. Table 1 reports localization uncertainty diagnostics for the high- and low-turbulence sub-volumes. Recall that $\tilde{\sigma}_z$ is the “prior center” for the depth-wise uncertainty, σ_z^{MAP} is the MAP estimate, σ_z^{res} is the residual-based estimate, and $e_{\sigma,z}$ is the relative error in σ_z^{res} with respect to the ground truth. The scale-balance factor $\kappa_z = (\sigma_z^{\text{res}}/\sigma_z)^2$ compares the residual-based estimate and the inferred/posterior value; κ_z^{MAP} is the fixed-point value predicted by Eq. (23); and $\kappa_z = \kappa_z^{\text{MAP}}$ indicates that the reconstruction has converged upon the MAP point and reached the desired prior–data balance. The true uncertainty for this case is $\sigma_z^* = 340 \text{ } \mu\text{m}$; cases where the prior center coincides with σ_z^* are indicated in Table 1 using bold text.

Table 1. Out-of-plane localization uncertainty diagnostics for the synthetic channel flow case. The prior center is defined as $\tilde{\sigma}_z = r_z \sigma_z^*$ and swept over $r_z \in \{0.25, 0.5, 1, 2, 4\}$, with $\sigma_z^* = 340 \text{ } \mu\text{m}$. Bold entries indicate the prior centered on σ_z^* .

Regime	$\tilde{\sigma}_z$ (μm)	σ_z^{MAP} (μm)	σ_z^{res} (μm)	$e_{\sigma,z}$ (%)	κ_z	κ_z^{MAP}
Low-turbulence	85	287.5	329.5	3.1	1.31	1.31
	170	306.8	329.8	3.0	1.16	1.15
	340	330.3	330.3	2.9	1.00	0.99
	680	359.3	330.8	2.7	0.85	0.84
	1360	396.5	331.0	2.6	0.70	0.69
High-turbulence	85	293.0	329.8	3.0	1.27	1.31
	170	312.5	330.0	2.9	1.12	1.15
	340	336.8	330.8	2.7	0.96	1.00
	680	366.3	331.0	2.6	0.82	0.85
	1360	405.5	333.5	1.9	0.68	0.70

Across both regimes and all prior choices, σ_z^{res} remains tightly clustered around $330 \text{ } \mu\text{m}$. The prescribed ground truth is $340 \text{ } \mu\text{m}$, and the relative error remains within 3%, even though the prior uncertainty varies from $\tilde{\sigma}_z = 85$ to $1360 \text{ } \mu\text{m}$. The MAP estimate behaves differently: it increases from about $\sigma_z^{\text{MAP}} = 290$ to $400 \text{ } \mu\text{m}$ across the sweep. Hence, σ_z consistently updates toward σ^* but

does not necessarily converge to it. Of course, this trend is expected. When the prior is centered far from the true value, the data shift the estimate toward the truth; however, the prior still exerts some influence on the MAP estimate (*as it should!*). By contrast, σ_z^{res} is computed directly from the residuals between the latent and detected particle positions and thus remains close to the true level of noise so long as the reconstruction is accurate. This result is likely to depend upon the true level of uncertainty, the seeding density, and the number of tracks. As of writing this text, however, we have yet to assess those effects.

The scale-balance factor κ_z agrees closely with the fixed-point prediction κ_z^{MAP} in all ten reconstructions. Deviations between κ_z and κ_z^{MAP} are slightly larger for the high-turbulence sub-volume, which is consistent with the larger reconstruction errors reported in the next section. That is, real fluctuations in $\mathbf{u}$ which are not resolved in the reconstruction can also contribute to the particle residuals; as such, σ_z^{res} depends on both the localization error and reconstruction error. Even so, it remains within 3% of the true value across the present tests, indicating that σ_z^{res} is fairly robust to prior mis-specification for the LPT conditions simulated in this section.

5.2.2. Particle and flow field reconstruction

Next, we examine the robustness (or lack thereof) of reconstructed particle position and flow fields to the assumed depth-wise uncertainty. Figure 7 compares fixed- σ and variable- σ results as $r_z = \tilde{\sigma}_z/\sigma_z^*$ is swept from 0.25 to 4 for both the low- and high-turbulence sub-volumes. From left to right, the panels show the z -direction latent particle position error e_z , the wall-normal velocity error e_w , and the pressure error e_p . Errors in the latent particle positions are normalized by the friction length, and flow field errors are volume-averaged NRMSEs of the fluctuating components. Wall-normal quantities are reported because they are the most sensitive to the assumed depth-wise localization uncertainty. Representative flow fields are plotted below in Fig. 8, and the complete set of particle position and flow field errors is reported at the end of this section in Table 2.

Remarkably, the variable- σ reconstructions are nearly insensitive to $\tilde{\sigma}_z$, i.e., all of the dashed lines are essentially flat. In the low-turbulence case, the latent position error, wall-normal velocity error, and pressure error stay near $e_z = 0.31$, $e_w = 0.17$, and $e_p = 1.4$. In the high-turbulence regime, these same errors stay near $e_z = 0.42$, $e_w = 0.36$, and $e_p = 0.38$. Particle position errors in the z -direction are well below the raw localization uncertainty, which is $\sigma_z^*/\ell_\nu = 1.36$, showing that the optimized tracks are smoother (and more accurate) than the noisy LPT data. Greater velocity error in the near-wall, high-turbulence sub-volume is expected because it contains stronger fluctuations and smaller-scale motions, which are more difficult to distinguish from measurement noise. Elevated pressure errors in the low-turbulence case are mainly an effect of the normalization. That is, pressure fluctuations are weaker further away from the wall, so the denominator based on p' is smaller for the low-turbulence domain and the relative error is amplified, accordingly.




Figure 7. Reconstruction errors for the channel flow case as $r_z = \tilde{\sigma}_z / \sigma_z^*$ is varied. Fixed- σ reconstructions prescribe $\sigma_z = \tilde{\sigma}_z$, while variable- σ reconstructions infer σ_z during training. Left: z -direction position error. Middle: w -velocity error. Right: pressure error. Colors distinguish the low- and high-turbulence sub-volumes, and the vertical dotted lines mark $r_z = 1$.

A striking feature of these results is that the variable- σ reconstructions outperform the fixed- σ reconstructions even at $r_z = 1$, at which point the prescribed value of σ_z equals with its true value. Here, the variable- σ reconstruction still yields lower particle position errors, velocity errors, and pressure errors in both turbulence regimes. This can be seen in Fig. 7, where all the dashed lines fall below the corresponding solid lines, and in Table 2. The improvements at $r_z = 1$ are modest, to be sure, but we believe this trend lends further support to our proposed method. We do not yet have a complete explanation for this behavior, although one possibility is that the inferred uncertainty can partially compensate for model errors that are not explicitly represented in the likelihood formulation.

Although the fixed- σ reconstructions perform well at $r_z = 1$, they are much more sensitive to the prescribed uncertainty, $\tilde{\sigma}_z$, than the variable- σ reconstructions. When σ_z is underestimated, the optimizer treats noise-induced oscillations in x_z as if they were caused by bona fide fluid motions. The reconstruction is then overfitted to these noisy data, increasing errors in both the latent particle positions and the estimated flow fields. This failure mode is most apparent in the low-turbulence case, where the pressure fluctuations are weak and easily contaminated by erroneous velocity fluctuations. At $r_z = 0.25$, the fixed- σ reconstruction gives $e_z = 0.627$, which is about twice the variable- σ value of $e_z = 0.307$. The pressure error increases from $e_p = 1.379$ to 7.940, or about *six times* the variable- σ value, and the Pearson correlation drops from 0.89 to 0.37. The wall-normal velocity error also increases from 0.166 for the variable- σ reconstruction to 0.413 for the fixed- σ reconstruction. The same trend appears in the high-turbulence cases, albeit with smaller relative differences, likely due to the lower signal-to-noise ratio and higher overall errors.

Overestimating σ_z produces the opposite failure mode. The constraint imposed by the z -component track data becomes too weak, and the reconstruction relaxes toward an overly smooth solution. The adverse effects caused by overestimating σ_z are milder than those caused by underestimating it, but the latent particle position and velocity errors remain above the variable- σ levels. This fail-

ure mode is exacerbated in the high-turbulence case, wherein the data contain more small-scale motions. At $r_z = 4$, for instance, the fixed- σ reconstruction gives $e_z = 0.757$ and $e_w = 0.461$, compared with $e_z = 0.436$ and $e_w = 0.366$ for the variable- σ reconstruction.




Figure 8. Representative w -velocity reconstructions for the low-turbulence sub-volume of the channel flow case at $r_z = 0.25$ and $r_z = 4$. Fixed- σ reconstructions are sensitive to the prescribed depth-wise uncertainty; variable- σ reconstructions remain stable across the sweep.

The trends discussed thus far can be seen qualitatively in Fig. 8, which depicts w -velocity reconstructions from the low-turbulence sub-volume. The top row shows fixed- σ results for $r_z = 0.25$ and $r_z = 4$, and the bottom row shows the variable- σ counterparts. Reconstructed fields are shown on the left side and error magnitude fields are shown on the right side; the ground truth reference is the low-turbulence panel in Fig. 5. Fixed- σ reconstructions exhibit a noticeable dependence on r_z , with spurious high-frequency features at $r_z = 0.25$, i.e., when the prior is too tight and the reconstruction is overfitted to noise. These features are suppressed at $r_z = 4$, but the corresponding velocity field is overly smooth. By contrast, the variable- σ fields remain stable across the sweep and have lower reconstruction error than their fixed- σ counterparts in all cases.

These results collectively demonstrate that fixed- σ reconstruction performs well only when the prescribed value of σ_z is accurate: a condition that does not hold in many experiments. Indeed, in the upcoming sections, we assess reconstructions of experimental LPT data based on standard estimates of σ_z . In the synthetic test, we find that variable- σ formulation mitigates the dependence of reconstructions on the assumed uncertainty by inferring it jointly with the particle trajectories and the flow field. This approach seems to yield accurate estimates of the depth-wise uncertainty while also producing particle, velocity, and pressure reconstructions that remain accurate in spite of large errors in $\tilde{\sigma}_z$.

Table 2. Reconstruction errors for the synthetic channel flow case. Particle position errors are normalized by the friction length, $\ell_\nu = 250 \mu\text{m}$, and flow field errors are NRMSEs of the fluctuating components averaged over the measurement volume. Bold entries indicate the prior centered on σ_z^* .

Regime	Mode	$\tilde{\sigma}_z$ (μm)	Particle position			Flow field			
			e_x	e_y	e_z	e_u	e_v	e_w	e_p
Low-turbulence	Variable- σ	85	0.046	0.046	0.307	0.131	0.101	0.166	1.379
		170	0.047	0.046	0.307	0.134	0.102	0.167	1.418
		340	0.047	0.046	0.307	0.138	0.105	0.166	1.362
		680	0.046	0.046	0.308	0.131	0.102	0.166	1.378
		1360	0.046	0.046	0.309	0.134	0.102	0.166	1.388
	Fixed- σ	85	0.062	0.063	0.627	0.318	0.256	0.413	7.940
		170	0.055	0.056	0.401	0.227	0.184	0.291	3.851
		340	0.048	0.048	0.334	0.155	0.124	0.191	1.637
		680	0.049	0.048	0.431	0.164	0.130	0.247	1.624
		1360	0.049	0.048	0.363	0.160	0.125	0.244	1.637
High-turbulence	Variable- σ	85	0.050	0.050	0.408	0.369	0.304	0.361	0.379
		170	0.050	0.050	0.408	0.365	0.301	0.354	0.380
		340	0.050	0.050	0.415	0.365	0.303	0.359	0.376
		680	0.050	0.050	0.417	0.365	0.301	0.357	0.381
		1360	0.050	0.050	0.436	0.368	0.305	0.366	0.382
	Fixed- σ	85	0.071	0.074	0.624	0.385	0.318	0.358	0.677
		170	0.067	0.068	0.500	0.370	0.303	0.356	0.447
		340	0.066	0.067	0.482	0.375	0.312	0.375	0.406
		680	0.067	0.069	0.792	0.393	0.328	0.450	0.440
		1360	0.067	0.069	0.757	0.396	0.331	0.461	0.443

5.3. Application to an experimental boundary layer dataset

We next apply the method to the experimental ZPG-TBL dataset described in Sec. 3.2. Since no ground truth is available, our analysis focuses on two checks. First, we compare fixed- σ and variable- σ reconstructions across a sweep of prescribed or prior depth-wise uncertainties. Second, we use the diagnostic relation $\kappa_z = \kappa_z^{\text{MAP}}$ to assess whether the variable- σ reconstructions have reached the prior–data balance implied by the MAP objective. The initial sweep is centered on the TrackFit estimate, $\sigma_z^* = 12 \mu\text{m}$, obtained from DaVis 10.2.1. The depth-wise uncertainty prior is swept according to $r_z = \tilde{\sigma}_z/\sigma_z^* \in \{0.25, 0.5, 1, 2, 4\}$ using this reference.

Table 3 reports the inferred localization uncertainties. The prior center is increased from $\tilde{\sigma}_z = 3$ to $48 \mu\text{m}$, all the while the residual scale remains between $\sigma_z^{\text{res}} = 14.8$ and $15.4 \mu\text{m}$. These estimates

place the depth-wise uncertainty near $15\text{ }\mu\text{m}$, about 25% larger than the TrackFit estimate used to center the initial sweep. As in the synthetic case, the MAP estimate depends on the prior center, growing from 12.4 to $17.8\text{ }\mu\text{m}$ as $\tilde{\sigma}_z$ is increased. By contrast, σ_z^{res} is nearly invariant across the sweep. The measured κ_z also agrees well with κ_z^{MAP} for all runs, indicating that the reconstructions have converged upon the MAP point. These trends mirror the behavior observed in the synthetic case and thus support the use of our variable- σ framework for this dataset. Moreover, the simple geometric estimate $\sigma_z = \sigma_{xy} / \sin(\theta/2)$ gives $\sigma_z \approx 15\text{ }\mu\text{m}$ for this imaging configuration, which is in close agreement with the residual-based estimates in Table 3. This agreement provides evidence that the Bayesian estimate of σ_z is more accurate than the TrackFit estimate for this dataset. That is not necessarily surprising, however, since the current method incorporates a physics-constrained estimate that incorporates the full flow field, whereas TrackFit is based on the track statistics, alone.

Table 3. Out-of-plane localization uncertainty diagnostics for the experimental ZPG-TBL case. The prior center is defined as $\tilde{\sigma}_z = r_z \sigma_z^*$ and swept over $r_z \in \{0.25, 0.5, 1, 2, 4\}$, with $\sigma_z^* = 12\text{ }\mu\text{m}$ from TrackFit. Bold entries indicate the prior centered on σ_z^* .

$\tilde{\sigma}_z$ (μm)	σ_z^{MAP} (μm)	σ_z^{res} (μm)	κ_z	κ_z^{MAP}
3	12.4	14.8	1.42	1.35
6	13.7	15.0	1.20	1.21
12	14.9	15.2	1.05	1.05
24	15.9	15.1	0.90	0.90
48	17.8	15.4	0.75	0.75

Figure 9 shows representative TBL reconstructions. The lower plane is colored by wall shear stress (WSS), and isosurfaces of the Q -criterion field, which represent vortical structures, are superimposed on the volume. The isosurfaces are colored by the wall-normal velocity. As before, the top row corresponds to fixed- σ reconstructions and the bottom row to variable- σ reconstructions. From left to right, the columns show results for $r_z = 0.25, 1$, and 4 , relative to the TrackFit reference value of $\sigma_z^* = 12\text{ }\mu\text{m}$. The vortical structures in the fixed- σ reconstructions change noticeably with the assumed wall-normal uncertainty, and the WSS field also exhibits visible variations. At $r_z = 0.25$, the fixed- σ reconstruction contains more fragmented, small-scale Q -structures, especially near the wall. At $r_z = 4$, both the Q -structures and WSS field become smoother. By contrast, the variable- σ reconstructions are mutually consistent across the whole prior sweep, and they are also broadly consistent with the fixed- σ reconstruction at $r_z = 1$.

Following the above analysis, we performed one additional variable- σ reconstruction using an updated reference value of $\tilde{\sigma}_z = 15\text{ }\mu\text{m}$. This run gave $\kappa_z = 0.996$, in very close agreement with the fixed-point value of $\kappa_z^{\text{MAP}} = 0.997$. We therefore regard $15\text{ }\mu\text{m}$ as a reasonable estimate of the wall-normal localization uncertainty for this dataset.




Figure 9. Representative reconstructions of the experimental ZPG-TBL dataset. Rows show fixed- σ and variable- σ reconstructions; columns show $r_z = \tilde{\sigma}_z/\sigma_z^* = 0.25, 1$, and 4 , with $\sigma_z^* = 12 \mu\text{m}$. Vortical structures are shown using Q -criterion isosurfaces ($Q = 6 \times 10^5 \text{ s}^{-2}$) colored by w ; the lower plane is colored by WSS.

5.4. Application to an experimental fluid–structure interaction dataset

We finally apply the method to the experimental FSI dataset described in Sec. 3.3. Since no ground truth is available, the goal is again to assess whether the inferred uncertainty is stable and whether the variable- σ reconstructions are resilient to the uncertainty prior center. For this case, TrackFit provided a depth-wise uncertainty estimate of $85 \mu\text{m}$, while the camera calibration yielded separation angles between $8\text{--}12^\circ$, corresponding to depth-wise uncertainties of approximately $162\text{--}244 \mu\text{m}$ via the geometric relationship presented in the introduction. For simplicity, we adopted an assumed uncertainty ratio of $\sigma_z/\sigma_{xy} = 10$, giving the reference value $\sigma_z^* = 170 \mu\text{m}$.

Table 4 reports the inferred out-of-plane localization uncertainty. The prior center is swept from $\tilde{\sigma}_z = 42.5$ to $680 \mu\text{m}$. Across these tests, the residual scale remains relatively stable, ranging from $\sigma_z^{\text{res}} = 238$ to $252 \mu\text{m}$. These values place the effective out-of-plane uncertainty near $245 \mu\text{m}$ for this experiment. This estimate is roughly three times larger than the TrackFit value and lies near the upper end of the range obtained from the simplified geometric estimate. The TrackFit and residual-based estimates correspond to σ_z/σ_{xy} ratios of approximately 5 and 14 , respectively, which is a conspicuous discrepancy. Further work is needed to corroborate this result and to identify the source of the discrepancy. As a preliminary consistency check, the measured values of κ_z agree with κ_z^{MAP} for all five tests, indicating converged reconstructions.

Figure 10 shows representative reconstructions of the w -velocity field for the FSI case. The top and bottom rows show the fixed- σ and variable- σ results, respectively. From left to right, the columns correspond to $r_z = 0.25, 1$, and 4 . The large-scale wake pattern is visually similar across

Table 4. Out-of-plane localization uncertainty diagnostics for the experimental FSI case. The prior center is defined as $\tilde{\sigma}_z = r_z \sigma_z^*$ and swept over $r_z \in \{0.25, 0.5, 1, 2, 4\}$, with $\sigma_z^* = 170 \mu\text{m}$ from the imaging-based reference estimate. Bold entries indicate the prior centered on σ_z^* .

$\tilde{\sigma}_z$ (μm)	σ_z^{MAP} (μm)	σ_z^{res} (μm)	κ_z	κ_z^{MAP}
42.5	202.1	237.9	1.39	1.39
85	217.5	241.4	1.23	1.23
170	235.8	244.8	1.08	1.08
340	257.9	248.4	0.93	0.93
680	285.5	252.0	0.78	0.78

all cases, but the small-scale structure of the w -velocity field changes with the prescribed uncertainty in the fixed- σ reconstructions. At $r_z = 0.25$, the fixed- σ velocity field contains sharper and more fragmented fluctuations, consistent with overfitting to noisy particle positions. At $r_z = 4$, the field becomes smoother, as in Secs. 5.2 and 5.3. By contrast, the variable- σ reconstructions are mutually consistent across the sweep, exhibiting similar wake structures and near-strut velocity patterns despite large changes to the prior. This stability provides qualitative support for the inferred uncertainty scale, since the reconstruction can absorb large changes in $\tilde{\sigma}_z$ without producing the *apparent* overfitting or oversmoothing observed in the fixed- σ cases.




Figure 10. Representative w -velocity reconstructions for the experimental FSI case. Rows show fixed- σ and variable- σ reconstructions; columns show $r_z = \tilde{\sigma}_z / \sigma_z^* = 0.25, 1$, and 4 , with $\sigma_z^* = 170 \mu\text{m}$. The strut cross-section is superimposed in black.

As an additional check, we performed a final variable- σ reconstruction using $\tilde{\sigma}_z = 240 \mu\text{m}$. This run gave $\kappa_z = 1.023$, in close agreement with $\kappa_z^{\text{MAP}} = 1.028$. Together with the stability of σ_z^{res} across

the prior sweep and the qualitative behavior in Fig. 10, this result supports an effective depth-wise uncertainty of approximately $240\text{ }\mu\text{m}$. Although no ground truth is available, the reconstruction evidence favors this value over the smaller TrackFit estimate, illustrating how Bayesian inference can supplement unreliable estimates of Σ_x .

6. Conclusions and outlook

This work presented a Bayesian extension of *neural-implicit particle advection* for joint particle–flow–uncertainty reconstruction from sparse and noisy LPT measurements. The approach treats the depth-wise localization uncertainty, σ_z , as an unknown nuisance parameter and estimates it together with the latent particle trajectories, velocity field, and pressure field. Measurement-based estimates of σ_z are used to define a weakly informative log-normal prior. The formulation adds a low-dimensional uncertainty variable and a prior term to the original NIPA objective loss; the computational cost is thus nearly identical to the fixed-uncertainty formulation.

The method was validated on a synthetic turbulent channel flow case with known ground truth and then applied to two experimental datasets: a zero-pressure-gradient turbulent boundary layer case and a fluid–structure interaction case involving a flexible strut in the wake of a cylindrical bluff body. In the synthetic case, the residual-based estimate of σ_z recovered the true value to within 3% across a sixteen-fold sweep of the prior center. In both experimental cases, for which there was no ground truth, the inferred depth-wise scale remained stable across the same prior sweep and satisfied our MAP fixed-point diagnostic, indicating convergence. Fixed- σ reconstructions were sensitive to the prescribed uncertainty: underestimating σ_z led to overfitting of noisy tracks and overestimating it weakened the data constraint and produced over-smooth fields. By inferring σ_z within the reconstruction, our variable- σ formulation reduced these sensitivities and produced more stable particle, velocity, and pressure estimates. Hence, the Bayesian formulation provides a practical way to calibrate σ_z in DA reconstructions of LPT data.

Several questions remain for future work. Additional tests are needed to assess the effects of seeding density, track length, noise level, and prior width on the inferred uncertainty. Future work should also extend the uncertainty model from a scalar unknown to the full tensor, Σ_x . Finally, we will extend framework to perform posterior UQ of the reconstructed velocity and pressure fields as well as downstream quantities of interest.

References

Bajpayee, A. (2014). *3D Particle Tracking Velocimetry Using Synthetic Aperture Imaging* (Unpublished doctoral dissertation). Massachusetts Institute of Technology.

- Barnkob, R., Kähler, C. J., & Rossi, M. (2015). General defocusing particle tracking. *Lab on a Chip*, 15(17), 3556–3560.
- Bhattacharya, S., & Vlachos, P. P. (2020). Volumetric particle tracking velocimetry (PTV) uncertainty quantification. *Experiments in Fluids*, 61(9), 197.
- Di Napoli, I. M., & Harwood, C. M. (2020). Design and benchmarking of a robust strain-based 3D shape sensing system. *Ocean Engineering*, 201, 107071.
- Fuchs, T., Bross, M., & Kähler, C. J. (2023). Wall-shear-stress measurements using volumetric μ PTV. *Experiments in Fluids*, 64(6), 115.
- Gesemann, S. (2021). TrackFit: uncertainty quantification, optimal filtering and interpolation of tracks for time-resolved Lagrangian particle tracking. In *14th International Symposium on Particle Image Velocimetry*. Chicago, Illinois.
- 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* (Vol. 4). Lisbon, Portugal.
- Hall, E. M., Guildenbecher, D. R., & Thurow, B. S. (2017). Uncertainty characterization of particle location from refocused plenoptic images. *Optics Express*, 25(18), 21801–21814.
- Hans, A., Bhattacharya, S., Hao, K., Vlachos, P., & Bilionis, I. (2024). Bayesian reconstruction of 3D particle positions in high-seeding density flows. *Measurement Science and Technology*, 35(11), 116002.
- He, C., Li, S., & Liu, Y. (2025). Data assimilation: new impetus in experimental fluid dynamics. *Experiments in Fluids*, 66(5), 1–24.
- Janke, T., & Michaelis, D. (2021). Uncertainty quantification for PTV/LPT data and adaptive track filtering. In *14th International Symposium on Particle Image Velocimetry*. Chicago, Illinois.
- Jux, C., Sciacchitano, A., & Scarano, F. (2020). Flow pressure evaluation on generic surfaces by robotic volumetric PTV. *Measurement Science and Technology*, 31(10), 104001.
- Kulkarni, Y., Stansberry, D., Di Napoli, I., Buchholz, J., & Harwood, C. (2024). A co-analysis framework for the visualization of synchronous fluid–structure modes in experimental FSI. In *Proceedings of the 35th Symposium on Naval Hydrodynamics*. Nantes, France.
- Kumar, S., & Hong, J. (2025). A review of 3D particle tracking and flow diagnostics using digital holography. *Measurement Science and Technology*, 36(3), 032005.

- Maas, H.-G., Gruen, A., & Papantoniou, D. (1993). Particle tracking velocimetry in three-dimensional flows: Part 1. Photogrammetric determination of particle coordinates. *Experiments in fluids*, 15(2), 133–146.
- Moaven, M., Gururaj, A., Raghav, V., & Thurow, B. (2024). Improving depth uncertainty in plenoptic camera-based velocimetry. *Experiments in Fluids*, 65(4), 49.
- Perlman, E., Burns, R., Li, Y., & Meneveau, C. (2007). Data exploration of turbulence simulations using a database cluster. In *Proceedings of the 2007 ACM/IEEE Conference on Supercomputing* (p. 23). Reno, Nevada.
- Rajendran, L. K., Bhattacharya, S., Bane, S. P., & Vlachos, P. P. (2021). Meta-uncertainty for particle image velocimetry. *Measurement Science and Technology*, 32(10), 104002.
- Rival, D. E., & van Oudheusden, B. (2017). Load-estimation techniques for unsteady incompressible flows. *Experiments in Fluids*, 58(3), 20.
- Schanz, D., Schröder, A., Gesemann, S., Huhn, F., Novara, M., Geisler, R., ... Depuru-Mohan, K. (2018). Recent advances in volumetric flow measurements: high-density particle tracking ('Shake-The-Box') with Navier–Stokes regularized interpolation ('FlowFit'). *New Results in Numerical and Experimental Fluid Mechanics XI*, 587–597.
- Schröder, A., Geisler, R., Sieverling, A., Wieneke, B., Henning, A., Scarano, F., ... Poelma, C. (2009). Lagrangian aspects of coherent structures in a turbulent boundary layer flow using TR-Tomo PIV and PTV. In *8th International Symposium on Particle Image Velocimetry* (pp. 25–28). Melbourne, Australia.
- Schröder, A., & Schanz, D. (2023). 3D Lagrangian particle tracking in fluid mechanics. *Annual Review of Fluid Mechanics*, 55, 511–540.
- Sciacchitano, A. (2019). Uncertainty quantification in particle image velocimetry. *Measurement Science and Technology*, 30(9), 092001.
- Sciacchitano, A., & Wieneke, B. (2016). PIV uncertainty propagation. *Measurement Science and Technology*, 27(8), 084006.
- Slotnick, J. P., Khodadoust, A., Alonso, J., Darmofal, D., Gropp, W., Lurie, E., & Mavriplis, D. J. (2014). *CFD Vision 2030 Study: A Path to Revolutionary Computational Aerosciences* (Tech. Rep. No. NF1676L-18332). Hampton, VA: National Aeronautics and Space Administration.
- Tang, R., Zhou, K., Tan, J., & Grauer, S. J. (2025). Neural inference of fluid–structure interactions from sparse off-body measurements. *arXiv preprint arXiv:2506.23480*.

- Toloui, M., & Hong, J. (2015). High fidelity digital inline holographic method for 3D flow measurements. *Optics Express*, 23(21), 27159–27173.
- 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.
- Wang, S., Li, B., Chen, Y., & Perdikaris, P. (2024). PirateNets: Physics-informed deep learning with residual adaptive networks. *Journal of Machine Learning Research*, 25(402), 1–51.
- Wang, S., Teng, Y., & Perdikaris, P. (2021). Understanding and mitigating gradient flow pathologies in physics-informed neural networks. *SIAM Journal of Scientific Computing*, 43(5), A3055–A3081.
- Wernet, M. P., & Pline, A. (1993). Particle displacement tracking technique and Cramer-Rao lower bound error in centroid estimates from CCD imagery. *Experiments in Fluids*, 15(4), 295–307.
- Westerweel, J. (2000). Theoretical analysis of the measurement precision in particle image velocimetry. *Experiments in Fluids*, 29(Suppl 1), S003–S012.
- Wieneke, B. (2013). Iterative reconstruction of volumetric particle distribution. *Measurement Science and Technology*, 24(2), 024008.
- Yamakaitis, M. J., Pigeon, A. M., Ehrmann, R., Parker, C. M., Anderson, J., & Bardet, P. M. (2025). Lagrangian particle tracking with a single camera in the viscous layer and buffer region of a zero-pressure-gradient turbulent boundary layer. In *16th International Symposium on Particle Image Velocimetry*. Tokyo, Japan.
- Yamakaitis, M. J., Rezaigui, A., Fort, C., Capanna, R., & Bardet, P. M. (2025). Four-view single-camera plenoptic 3.0 imager for three-dimensional Lagrangian particle tracking in an oscillatory laminar boundary layer. *Experiments in Fluids*, 66(9), 165.
- Young, Y. L., Valles, Z., Di Napoli, I., Montero, F. M., Minerva, L. F., & Harwood, C. (2023). Wave effects on the hydroelastic response of a surface-piercing hydrofoil. Part 1. Fully wetted and ventilated flows. *Journal of Fluid Mechanics*, 963, A37.
- Zaki, T. A., & Wang, M. (2025). Data assimilation and flow estimation. In *Data Driven Analysis and Modeling of Turbulent Flows* (pp. 129–181). Elsevier.
- Zhang, J., Bhattacharya, S., & Vlachos, P. P. (2020). Using uncertainty to improve pressure field reconstruction from PIV/PTV flow measurements. *Experiments in Fluids*, 61(6), 131.
- Zhou, K., & Grauer, S. J. (2026). On the joint estimation of flow fields and particle properties from Lagrangian data. *arXiv preprint arXiv:2510.00479*.