Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Accurate estimation of velocity and acceleration, their statistics and position prediction in Lagrangian Particle Tracking

Tom Buchwald

buchwald@b-tu.de

Brandenburg Technical University Cottbus-Senftenberg

Daniel Schanz

German Aerospace Center (DLR)

Philipp Godbersen

German Aerospace Center (DLR)

Andreas Schröder

Brandenburg Technical University Cottbus-Senftenberg

Research Article

Keywords: Regularized B-splines, Lagrangian Particle Tracking, Signal processing, Denoising

Posted Date: July 17th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10342778/v1

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

Additional Declarations: No competing interests reported.

Accurate estimation of velocity and acceleration, their statistics and position prediction in Lagrangian particle tracking

Tom Buchwald[1*], Daniel Schanz[2], Philipp Godbersen[2], Andreas Schröder[1,2]

[1]Chair of Image based measurement techniques, Brandenburg Technical University Cottbus-Senftenberg, Siemens-Halske-Ring 15a, Cottbus, 03046, Brandenburg, Germany.
[2]Institute of Aerodynamics and Flow Technology, Experimental Methods, German Aerospace Center (DLR), Bunsenstraße 10, Göttingen, 37073, Niedersachsen, Germany.

*Corresponding author(s). E-mail(s): buchwald@b-tu.de;
Contributing authors: Daniel.Schanz@dlr.de;
Philipp.Godbersen@dlr.de; Andreas.Schroeder@dlr.de;

Abstract

This study presents complementary approaches to mitigate the impact of random measurement noise on velocity and acceleration statistics obtained from Lagrangian Particle Tracking (LPT). Using synthetically generated particle tracks in a turbulent near-wall channel flow, we assess both track-level processing and statistical noise separation strategies.
First, an adaptive extension of the conventional TrackFit algorithm is introduced, in which filter parameters are adjusted to local flow conditions rather than prescribed globally. This local TrackFit approach enables accurate estimation of instantaneous velocities and, in particular, accelerations, even in strongly inhomogeneous regions such as the viscous sublayer, where conventional global filtering leads to significant bias. Furthermore, this extension of the TrackFit concept enables more accurate prediction of particle positions, as required in Shake-The-Box (STB) algorithms.
Second, track-merging strategies are assessed to combine statistically independent track segments of the same particle obtained from multiple temporally shifted tracking passes, thereby enhancing accuracy in variable-timestep LPT.

Third, a statistical noise correction framework based on Gaussian error propagation is introduced, in which measurement noise is explicitly separated from the second-order moment estimates of velocity and acceleration. This approach does not rely on strong temporal filtering and therefore avoids underestimation of the corresponding moments. Together, these methods enable accurate estimation of instantaneous Lagrangian velocities and accelerations through adaptive track filtering, while allowing noise-robust velocity and acceleration statistics to be obtained without filtering via explicit noise subtraction.

Keywords: Regularized B-splines , Lagrangian Particle Tracking, Signal processing, Denoising

1 Introduction

A significant obstacle to achieving a fundamental understanding of turbulent flows and developing practical, sufficiently accurate simulation models for engineering problems is a lack of precise measurement data. This is particularly true for acceleration data, which leaves a gap in turbulence research regarding the theoretical prediction of acceleration statistics, despite its practical relevance. This is exemplified by the Heisenberg–Yaglom relation (Monin and Yaglom 1975) and the velocity-acceleration structure function (Mann et al. 1999). Consequently, the validity of such theoretical predictions is only ascertained in exceptional cases.

Despite the advent of highly sophisticated LPT algorithms capable of directly estimating particle acceleration, the limited positional precision of particle tracks still leads to inaccuracies in acceleration measurements due to the required double differentiation. As a result, acceleration data are typically incorporated into reconstruction algorithms such as FlowFit3 (Godbersen et al. 2024), but only with comparatively low weighting in the cost function.

Not only more accurate acceleration, but also precise position and velocity measurements are increasingly demanded, particularly as input for flow reconstruction with physics-informed neural networks (PINNs) (Raissi et al. 2019). PINNs have demonstrated superior accuracy over conventional methods, particularly in low seeding density measurements (Shin and Schröder 2025), and can even account for particle inertia effects (Zhou and Grauer 2023). Furthermore, the latter approach can filter noise from raw particle tracks together with the flow reconstruction.

Unlike standard filtering techniques, this method leverages Navier-Stokes constraints to account for the entire flow field globally. While this yields highly convincing results, particularly at high seeding densities (Zhou and Grauer 2026), it requires significant computational effort. To the authors' knowledge, the extent to which a prior filtering stage can further improve these results has not yet been investigated. However, it is well established for other PINN architectures that, while they are robust to moderate noise (DiLeoni et al. 2023), increased noise degrades reconstruction quality.

Measurement noise in LPT can arise from multiple sources, including camera sensor noise, limited temporal sampling, and particle image overlap.

To reduce its impact on derived quantities, contemporary LPT techniques commonly apply track-based filtering approaches to smooth particle tracks. Unlike the PINN-based approach, these methods are computationally inexpensive because they operate on each track independently, processing individual tracks without considering the global constraints of the surrounding flow field.

Techniques based on polynomial regression, such as the Savitzky-Golay filter (Savitzky and Golay 1964), are commonly employed for this purpose. These filters convolve a polynomial of a specified degree with the measured signal to produce a smoothed track. Alternatively, convolution with a Gaussian kernel can be used (Mordant et al. 2004). Another approach, TrackFit (Gesemann et al. 2016), represents tracks using B-splines and approximates an optimal Wiener filter (Wiener 1949) by solving an overdetermined linear system derived from the positional data and appropriate weights. TrackFit significantly outperforms both Savitzky-Golay filters (Gesemann 2021; Schmid et al. 2022; Lawson et al. 2018) and Gaussian kernel-based filters (Kearney et al. 2024) in high-frequency noise suppression, which is critical for accurate velocity and, in particular, acceleration measurements.

Recently, a Maximum Likelihood Estimator (MLE) filter (Kearney et al. 2024) was introduced, demonstrating performance comparable to that of TrackFit. A subsequent extension of this method, which incorporates intermittency effects, has even outperformed TrackFit in synthetic homogeneous isotropic turbulence test cases (Kearney et al. 2026). It is evident that all of these filters depend on effective parameterization to achieve optimal noise suppression. In the context of polynomial filters, the 'gut feeling' of experts is frequently employed. In the case of MLE filters, Kearney et al. employed an exhaustive parameter search using ground truth data. In contrast, the TrackFit approach determines optimal filter parameters directly from noisy experimental data by analyzing the amplitude spectral density (ASD) of the reconstructed tracks. The ASD also provides an accurate estimation of the standard deviation of position measurements, which facilitates the enhancement of post-processing steps, such as pressure reconstruction (Zhang et al. 2020).

Although optimal filter parameters evidently depend on the flow conditions at hand, previous filtering approaches have typically employed constant filter parameters for each particle track. The work of (Cheminet et al. 2021) demonstrated that selecting appropriate filter parameters inherently involves a trade-off between accurately reconstructing low- and high-acceleration flow regions. Consequently, adapting the regularization strength to local flow properties could provide substantial benefits for turbulence measurements, particularly when investigating phenomena such as intermittency. This need has become even more pronounced with the advent of advanced LPT techniques, such as Variable-Timestep Shake-the-Box (VT-STB) (Schanz et al. 2021), which enable precise measurements in highly inhomogeneous flows.

This development underscores the need for track filtering methods that can effectively suppress noise in highly variable flows and regions of minimal particle displacement. To date, only two locally adjusted, flow-agnostic track filtering approaches have been developed. The first is based on polynomial filtering, in which the filter length is determined from the standard deviation of particle positions estimated via

polynomial regression (Janke and Michaelis 2021). The second is a locally adapted TrackFit approach presented in (Buchwald et al. 2025).

While both approaches enable local adaptation of the track-fitting process, they exhibit important limitations. The polynomial-filter-based method shows reduced accuracy compared to the locally adapted TrackFit approach. The latter is currently restricted to stationary flow conditions, limiting its applicability to a narrower range of experimental configurations.

In this work, an improved locally and temporally adapted TrackFit approach is introduced, aiming to provide increased accuracy in instantaneous measurements while remaining applicable to general, non-stationary flows. As TrackFit is based on Wiener filter theory, which can be extended to predict particle positions as required in LPT algorithms such as STB (Schanz et al. 2016), the adaptive extension presented here also offers potential as an efficient predictor.

Furthermore, algorithms are proposed that merge independent, temporally shifted tracks to enable accurate LPT in datasets with small particle displacements.

In addition to improving instantaneous measurements, determining time averaged statistics is also of great interest. A simple yet effective method of removing random noise from data involves the implementation of two independent measurements facilitated by an auxiliary LPT system (Lawson et al. 2018).

However, with LPT systems becoming increasingly complex to achieve precise volumetric measurements, the use of two independent systems has grown impractical. Consequently, this study introduces and evaluates another method, which leverages the estimated standard deviation of noise in position measurements and employs Gaussian error propagation to compensate for the contribution of white noise to second-order moment estimates of velocity and acceleration.

The article is structured as follows: Section 2.1 introduces the synthetic test data sets. A description of the aforementioned methods is provided in Section 3. This is followed by an evaluation of the performance of the described approaches in Section 4. Section 5 then provides a brief assessment of the potential contribution of the presented methods to future developments.

2 Synthetic test cases

2.1 Channel flow

We analyzed approximately 10^6 simulated fluid particle tracks over 300 time steps in a turbulent channel flow from viscous to the end of the buffer layer $1 < y+ < 30$ at $Re_\tau = 1000$. The direct numerical simulation (DNS) data was obtained from Johns Hopkins Turbulence Database (Perlman et al. 2007; Li et al. 2008). Particles were initially seeded uniformly throughout the flow. During the simulation, additional particles were continuously introduced at the edges of the simulated volume to ensure that new particles entered the domain as others exited.

A dimensionless sampling interval of $h_s = 0.013$, corresponding to twice the DNS time step size, was chosen for particle advection. At each sampling time, flow velocities were obtained by linear interpolation between the stored DNS velocity fields. The

advection was performed using a Runge-Kutta scheme with two substeps per advection step, resulting in a time step size of $h = 0.026$ for the tracks.

The Kolmogorov time scale τ_η was evaluated based on the dissipation profile from a turbulent channel flow with the same Re_τ (Lee and Moser 2015). The oversampling ratio, defined as the ratio of the Kolmogorov time scale to the particle tracking time step, was subsequently computed. This ratio ranges from 7.7 at $y^+ = 2$ to 13 at $y^+ = 30$, demonstrating that the chosen time step ensures the smallest turbulent scales are well resolved.

Although truncation of extreme acceleration events cannot be ruled out, these values still represent a realistic sampling for LPT experiments.

White noise with standard deviation $\sigma_n = 0.00045$ is superimposed on the tracks, which is 3% of the mean displacement. This level is comparable to noise amplitudes reported for well-calibrated experimental measurement systems and therefore provides a reasonable test case.

Ground-truth velocity and acceleration values were computed by utilizing a 5-stencil central differences scheme.

2.2 Cylinder wake flow

An additional dataset of approximately 200000 Lagrangian tracks, each containing up to 300 timesteps from the wake downstream of a smooth cylinder at a subcritical Reynolds number $Re = 3900$ (based on the diameter of the cylinder and the free-stream velocity) is used to asses the potential of a locally adapted Wiener predictor. The tracks correspond to the first subdomain in the dataset described in more detail in (Khojasteh et al. 2022).They are downsampled to every eighth timestep, limiting their maximum size to 37 and reducing the Kolmogorov time oversampling ratio from roughly 20 to 2.5. The position signal is scaled to achieve a realistic mean displacement of 5 pixels. Compared to the channel flow data, this dataset exhibits a distinct flow regime and a locally lower Kolmogorov time-oversampling ratio.

3 Methods

The following section introduces three approaches aimed at reducing measurement noise. First, an enhanced version of local TrackFit is presented. Second, we describe a track merging procedure that combines tracks of the same particle obtained from independent LPT passes. Third, we discuss a correction method that separates the noise contribution from measured central moments of velocity or acceleration.

The fundamental assumption underlying all three approaches is that the noise level can be determined from an analysis of the ASD. We employ an algorithm based on an autocorrelation of the position signal to obtain the ASD from raw particle tracks, as described in (Gesemann 2021). Previous experiments confirm that noise in STB-based measurements usually shows a flat spectrum. This can be attributed primarily to camera image noise, which is composed of Gaussian-distributed thermal noise from the underlying electronic circuitry and Poisson-distributed shot noise. The latter, given sufficient light intensity, can be well approximated by a Gaussian distribution. However, other error sources, such as the limited accuracy of the particle position

optimization algorithm (referred to as "Shaking" in STB (Schanz et al. 2016)), can introduce non-white noise components. The significance of this effect depends on several factors, including the mismatch between the optical transfer function used and the actual particle image, overlapping particles, and variations in particle image intensity caused by different camera perspectives.

An example of the resulting distribution of shaking noise in a Rayleigh-Bénard convection experiment (Bosbach et al. 2021) is presented in Fig. 1. The plot illustrates

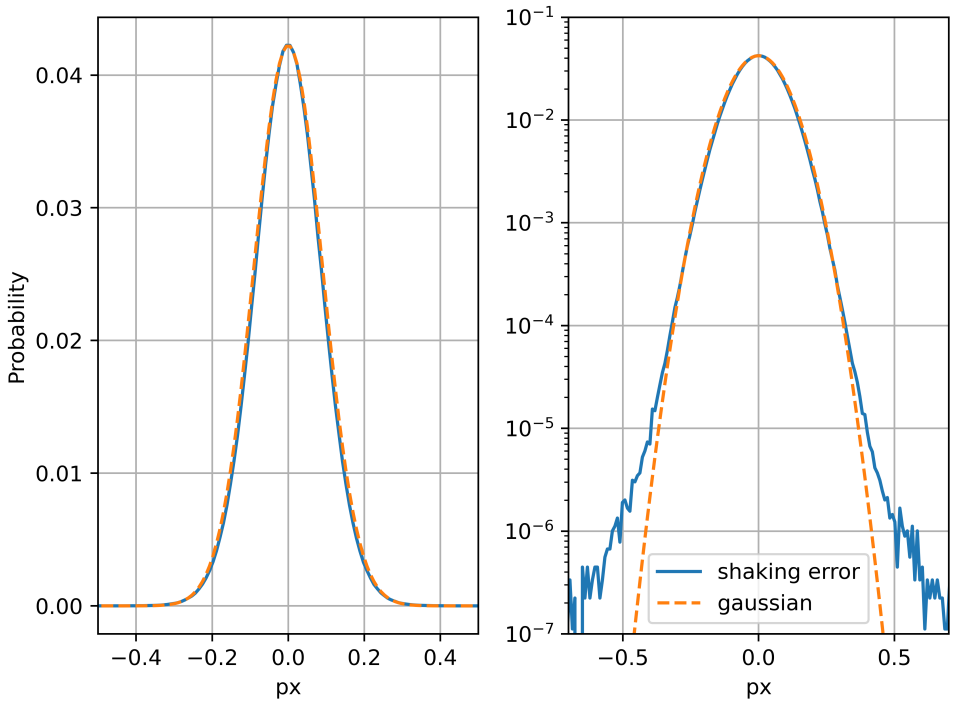


Fig. 1 Probability density function of the shaking error in a Rayleigh-Bénard convection Shake-The-Box experiment, compared with a normal distribution. Normal scaling (left) and semi-logarithmic scaling (right) are shown to highlight the deviation for extreme shaking errors

the deviation between raw particle positions and their fitted counterparts using Track-Fit. Since the fit attenuates only a fraction of the noise and a minor component of the true signal, the resulting distribution serves as an estimate of the actual shaking error only. Notably, the data suggest that the shaking noise in STB experiments exhibits a nearly Gaussian distribution, implying a flat spectrum.

Since this noise is highly dependent on the experimental setup and the specific implementation of the tracking algorithm, the resulting noise may exhibit different characteristics and does not always conform to a white spectrum, as demonstrated by (Cheminet et al. 2021).

A key distinction of the method presented here, compared to most commonly used filtering approaches, is the use of ASD evaluation to guide filter parameterization, thereby providing a framework for *a posteriori* uncertainty quantification. In many cases, parameters are chosen based on expert intuition (e.g., in polynomial regression), which does not guarantee an optimal choice, or fitted-data statistics are used as a quality criterion. (Küchler et al. 2024) demonstrated that, for a B-spline–based filter, a suitable temporal filter length can be identified by analyzing the acceleration variance across different lengths, determining the point where the signal transitions from noise-dominated to oversmoothed. However, this approach requires a large dataset, which hinders local parameter determination and is impractical for unsteady flows.

3.1 Local TrackFit

A brief description of the TrackFit methodology is provided here, with particular emphasis on the B-spline–based interpolation procedure. For a comprehensive presentation of the full TrackFit framework, the reader is referred to (Gesemann 2021).

To determine the optimal parameterization of the TrackFit filter, the ASD is computed from experimental track data, from which the signal-to-noise ratio (SNR) is obtained as a function of the normalized frequency f/f_{Nyquist}. The normalized frequency at which the SNR equals unity is defined as the cutoff frequency f_{cutoff}, which is used to configure the TrackFit filter. The choice $f_{\mathrm{cutoff}} = f_{\mathrm{SNR}=1}$ has been shown to effectively reduce the mean absolute error (MAE) compared to alternative cutoff frequency estimation methods (Cheminet et al. 2021).

However, this choice implies aggressive smoothing, which is unproblematic in terms of MAE for many tracks, but may lead to severe underestimation in high acceleration events. Consequently, no single SNR fits all purposes; the choice inherently depends on the chosen error metric. In particular, when the filtered quantities are intended as input for data assimilation methods like FlowFit, establishing the optimal balance between signal distortion and noise transmission becomes crucial. To date, a systematic investigation of this tradeoff remains an open gap in the literature and will be addressed in a forthcoming study.

In addition, the noise plateau in the high frequency range provides an estimate of the standard deviation of the white noise. This estimate is subsequently used to correct time-averaged quantities (see Section 3.3).

The TrackFit procedure is based on representing a particle track as a linear combination of cubic B-splines

$$B(t) = \begin{cases} \frac{1}{6}\,(2-|t|)^3, & 1 \le |t| \le 2, \\ \frac{1}{6}\,\left(4 - 6\,|t|^2 + 3\,|t|^3\right), & 0 \le |t| < 1, \\ 0, & |t| > 2. \end{cases} \tag{1}$$

and corresponding weigths $\vec{c} \in \mathbb{R}^{n+3}$.

Because the cubic B-spline is twice continuously differentiable, a track of length n with position $\vec{p} \in \mathbb{R}^n$, velocity $\vec{u} \in \mathbb{R}^n$ and acceleration $\vec{a} \in \mathbb{R}^n$ is continuously represented by the set of weights. The sampled quantities are obtained from the systems of equations

$$\vec{p} = E_0 \vec{c}, \qquad \vec{u} = -E_1 \vec{c}, \quad \vec{a} = E_2 \vec{c} \tag{2}$$

with a $\mathbb{R}^{n \times (n+3)}$ matrix

$$E_i = \frac{1}{h^i} \begin{pmatrix} \partial_t^i B(\tau-2) & \partial_t^i B(\tau-1) & \partial_t^i B(\tau) & \partial_t^i B(\tau+1) & & \\ & \partial_t^i B(\tau-2) & \partial_t^i B(\tau-1) & \partial_t^i B(\tau) & \partial_t^i B(\tau+1) & \\ & & \ddots & \ddots & \ddots & \ddots \\ & & & \partial_t^i B(\tau-2) & \partial_t^i B(\tau-1) & \partial_t^i B(\tau) & \partial_t^i B(\tau+1) \end{pmatrix} \tag{3}$$

for $0 \leq \tau < 1$. Here, $\tau = \Delta t / h$ denotes the temporal offset Δt of the sampled data relative to the acquisition time step size h. Each sampled value is thus obtained from the superposition of temporally shifted B-splines, each weighted by their corresponding coefficients.

The weights are determined by solving an additional linear system

$$\begin{pmatrix} E_0 \\ \frac{h^3}{w} E_3 \end{pmatrix} \vec{c} = \begin{pmatrix} \vec{s} \\ \vec{0} \end{pmatrix} \tag{4}$$

in a least squares sense. The parameter $w \approx (\pi \cdot f_{\text{cutoff}})^3$ controls the degree of smoothing applied to the track (Gesemann 2021), while $\vec{s}$ denotes a component of the raw position signal.

This type of noise reduction via discretely penalized B-splines (P-splines), first systematically introduced by (Eilers and Marx 1996), produces a fitted position signal $E_0 \vec{c}$ that closely resembles a signal obtained with an optimal Wiener filter, effectively minimizing the mean-squared position error. The design of the original TrackFit filter, however, relies on the implicit assumption that the SNR as a function of normalized frequency is identical for all tracks and at all times, allowing the use of a single cutoff frequency. As demonstrated in (Buchwald et al. 2025), this assumption does not hold in strongly inhomogeneous flows, such as round jets. In these cases, the filter may no longer be optimal, potentially introducing significant errors and even yielding nonphysical flow profiles.

In practice, the frequency at which noise begins to dominate the signal depends on the signal strength, that is, on the particle's velocity and acceleration. Accordingly, the spectral content of the ASD is linked to the turbulence intensity or Reynolds number, and thus to the local temporal scale separation. This requires that the parameter w be determined as a function of mean absolute velocity and acceleration.

To achieve this, a calibration of the mapping $f(\overline{|u_i|}, \overline{|a_i|}) = f_{\text{cutoff},i}$ is performed, with i representing the directional component. For simplicity, the subscript i is omitted for $f_{\text{cutoff},i}$ in the following. Tracks are segmented into sub-tracks of at most 50 time steps and pre-filtered using a global cutoff frequency. For each sub-track, the mean absolute velocity and acceleration are computed and used for a two-dimensional binning. The ASD is then calculated for each bin. An individual cutoff frequency

is determined as the frequency at which the SNR falls below one. Finally, bilinear interpolation of the resulting data provides a continuous f_{cutoff} field.

Filtering is subsequently conducted in two stages. First, a global filtering step is applied on the track. Then, for each timestep in the track, the local velocity and acceleration are evaluated. The mapping $f(\overline{|u_i|}, \overline{|a_i|})$ is used to determine the corresponding f_{cutoff}, and the associated parameters $\vec{w}$ are used to fit the raw track, such that Equation 4, with evaluated B-Splines for $\tau = 0$ is now reformulated to

$$\begin{pmatrix} \frac{1}{6} & \frac{2}{3} & \frac{1}{6} & & & & & \\ & \frac{1}{6} & \frac{2}{3} & \frac{1}{6} & & & & \\ & & \frac{1}{6} & \frac{2}{3} & \frac{1}{6} & & & \\ & & & \ddots & \ddots & \ddots & & \\ & & & & \frac{1}{6} & \frac{2}{3} & \frac{1}{6} & \\ -\frac{1}{w_1} & \frac{3}{w_1} & -\frac{3}{w_1} & \frac{1}{w_1} & & & & \\ & -\frac{1}{w_2} & \frac{3}{w_2} & -\frac{3}{w_2} & \frac{1}{w_2} & & & \\ & & -\frac{1}{w_3} & \frac{3}{w_3} & -\frac{3}{w_3} & \frac{1}{w_3} & & \\ & & & \ddots & \ddots & \ddots & \ddots & \\ & & & & -\frac{1}{w_n} & \frac{3}{w_n} & -\frac{3}{w_n} & \frac{1}{w_n} \end{pmatrix} \vec{c} = \begin{pmatrix} \vec{s} \\ \vec{0} \end{pmatrix} \tag{5}$$

with n denoting the number of timesteps in the track. Now, the set of weights $\vec{c}$ obtained by solving the shown equation system in a least squares sense corresponds to particle tracks which inherently adhere to the flow situation in each individual timestep. This represents a fundamental difference from the local TrackFit method introduced in (Buchwald et al. 2025), in which the filter parameter w depended on the spatial position of the track midpoint, whereas it is now determined by the velocity and acceleration estimated via a global TrackFit.

3.1.1 Local prediction

A local determination of the cutoff frequency f_{cutoff} is not only relevant for parameterizing a filter applied to fully reconstructed tracks. A key component of the STB algorithm relies on predicting particle positions at the subsequent time step. In our current STB implementation, this step employs a Wiener predictor parameterized by a global cutoff frequency (Schanz et al. 2016).

Using a locally determined f_{cutoff} parameterization for individual tracks promises increased prediction accuracy. While the mapping function for the channel flow dataset is two-dimensional, we employ a simplified one-dimensional mapping, $f(\overline{|a_i|}) = f_{\text{cutoff}}$, which excludes velocity because the cylinder wake dataset lacks a sufficient number of tracks for reliable 2D binning of ASDs.

To estimate acceleration, a globally parameterized track fit has been used on the known particle positions.

Instead of assigning a distinct f_{cutoff} to each timestep, as described for track filtering, we use a single cutoff frequency depending on mean absolute acceleration $\overline{|a_i|}$ to reduce implementation effort.

To maintain high prediction accuracy even in tracks with sudden acceleration changes in the most recent timesteps, $\overline{|a_i|}$ is computed as the linearly weighted mean of the absolute acceleration values over up to ten most recent timesteps, with the highest weight assigned to the latest timestep. The input is restricted to a maximum of 10 timesteps with linearly increasing weights, as this configuration previously yielded robust results in the global Wiener prediction. Furthermore, our implementation uses pre-computed Wiener coefficients for track lengths between 4 and 10 timesteps at f_{cutoff} increments of 0.05. Consequently, the computed f_{cutoff} for prediction must be rounded accordingly. In this work, we aim solely to introduce the concept of the local Wiener predictor and assess its potential. Therefore, a systematic investigation of implementation details, such as the number and weighting of input positions and accelerations, the quantization of f_{cutoff}, and the method for determining an optimal f_{cutoff} from the ASDs for prediction, is beyond the scope of this paper.

To evaluate the performance of the proposed local Wiener predictor, it is instructive to compare it against alternative predictive frameworks. Therefore, two baseline approaches are introduced for comparative testing:

1. **Quadratic Extrapolation:** A classic polynomial approach where the positions from the three preceding time steps are used to fit a second-order polynomial, which is then extrapolated to predict the position at the subsequent time step. This method serves as a well-established benchmark, as it is widely utilized within the particle tracking community.
2. **Taylor-Series Approximation:** As a simpler, physics-based alternative, a second-order Taylor-series expansion is evaluated, formulated as:

$$p_{n+1} = p_n + u_n \Delta t + \frac{1}{2} a_n \Delta t^2 \tag{6}$$

where p_n, u_n, and a_n represent a single component of position, velocity, and acceleration at the current time step, respectively, and Δt denotes the time step size. Although a viable alternative, this approach has rarely been the focus of methodological development in LPT. This lack of attention is likely due to its strong dependency on the accuracy of the velocity and acceleration estimates, meaning its success hinges entirely on robust preceding filtering. To assess this dependency, the Taylor-series predictor is evaluated using kinematic properties (u and a) derived from both global and local Trackfit formulations.

3.2 Track Merging for Individual Particles

STB distinguishes real from ghost particles based on the observation that real particles typically follow smooth tracks with small acceleration changes between time steps, whereas ghost particles appear at quasi-random locations. To initiate a new track, STB requires detections in multiple (typically four) consecutive frames with limited total acceleration, making the formation of ghost tracks unlikely under typical flow

conditions. However, this assumption only holds if the particle field exhibits sufficient relative motion between particles, such that ambiguities caused by a given particle image distribution decorrelate within a few time steps. While this assumption is valid for many flows, flows with a wide velocity range can contain regions with very little particle displacement from frame to frame. Jet flows, for example, typically exhibit a steep velocity gradient between the jet core and the surrounding flow. As the time separation between successive frames must be chosen to reliably track the fast particles within the jet (limiting their displacement to, e.g., 10–15 pixels), particles in the surrounding flow may move only a fraction of a pixel during the same interval. In such cases, ghost particles can persist across multiple frames and potentially form ghost tracks. Additionally, slowly evolving particle clouds increase the likelihood of persistent image overlap, thereby reducing particle position accuracy.

Therefore, achieving high quality STB tracking in flows with wide velocity range requires the application of variable timesteps (Schanz et al. 2021). In this approach, selected frames are deliberately skipped so that the particle displacement between successive evaluated images is sufficiently large to enable reliable tracking of the particles of a certain velocity regime. These tracks are subsequently interpolated onto a smaller timestep to generate particle positions for initializing the next STB pass (at a reduced timestep) for faster particles. This initialization with the positions of the slower particles is necessary because a high concentration of particles untrackable due to insufficient displacement impedes the identification of valid tracks. The VT-STB processing is not only important in flows with wide velocity ranges, but also serves to enhance the Kolmogorov time oversampling, thereby facilitating the measurement of extreme acceleration events.

The frame-skipping strategy not only improves trackability of particles with low displacement but also enables the generation of statistically independent track realizations. By shifting the evaluation sequence by at least one frame, the same particle can be tracked multiple times independently.

The question arises as to how the tracks from these datasets must be combined to the full time resolution to minimize noise. To this end, two different approaches are investigated in the following. First approach: The tracks representing the same particle across the different datasets are identified and their noisy position signals are combined. This allows the original timestep to be recovered. Subsequently, a local trackfit is applied. In the following, this procedure is termed as track interleaving. Second approach: The tracks representing the same particle across the different datasets are identified and interpolated onto the original timestep before being averaged. This averaging process reduces the noise as it is shown in Fig. 2. According to the $\sqrt{N}$ law, averaging N independent measurements decreases the standard deviation of uncorrelated noise by a factor of $\sqrt{N}$, thereby improving the SNR. Afterward, filtering can be applied if necessary. However, it should be noted that a local TrackFit is only optimal if the ASD matches the assumed model of particle motion. If the downsampling causes high frequency motion signals to exceed the Nyquist frequency, meaning the motion is severely undersampled, the aliasing effect will spectrally distort the signal and invalidate this assumption. Even if severe undersampling does not occur, the interpolation itself modifies the ASD, such that the applied filter is no longer optimal. This effect

cannot be attributed to suboptimal parameterization; instead, the interpolation alters the shape of the ASD itself, thereby reducing the separation between noise-dominated and signal-dominated frequency ranges.

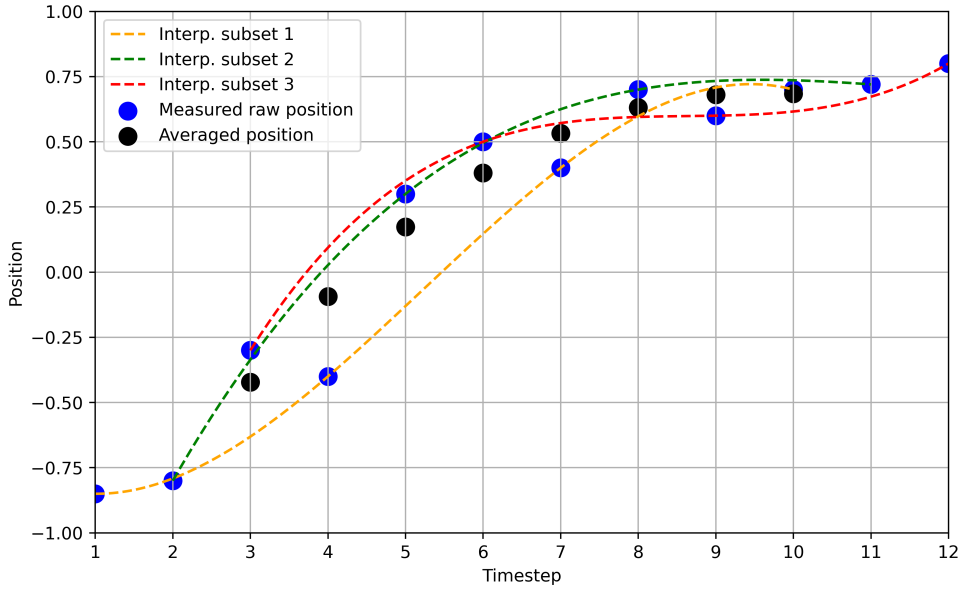


Fig. 2 Schematic of the track averaging principle for N=3 independently generated tracks with subsequent interpolation. Resulting mean values are shown as black dots, representing the consolidated position estimate

A critical step is the correct identification of corresponding interpolated tracks in the N data sets. In our implementation, a search radius proportional to the noise standard deviation is defined. For each track in the first interpolated dataset, the algorithm verifies whether a corresponding particle exists within this radius at all relevant time steps in the other dataset. A valid match is confirmed only if such a particle is present consistently on at least two time steps and is uniquely assigned to a single track.

Once corresponding tracks have been identified, their positions are either averaged or interleaved. In the case of track averaging, if a track is truncated in one dataset but continues in another, the averaging process can be sustained using fewer than N data points. Although this reduces accuracy, it enables the reconstruction of longer tracks. However, the results presented in this contribution are based on an implementation which considers only particles that were identified in all three datasets, in order to asses the best case noise reduction capability of the method. To suppress spurious associations - e.g. when particles that do not belong to the same track are erroneously linked - we impose an upper threshold on the standard deviation of the N averaged

positions. If this threshold is exceeded, the track is terminated a few time steps before the limit is reached.

3.3 Accurate time averaged statistics through noise subtraction

The finite accuracy of the position measurements introduces an additional contribution to the variance of time-derivative quantities such as velocity and acceleration. This contribution is independent of the flow and depends solely on the reconstruction procedure and the measurement uncertainty. Consequently, the second order moments obtained from the reconstructed signal comprises both a physical and a noise induced contribution. In the following, the latter is quantified.

Assuming Gaussian-distributed measurement noise that is uncorrelated with the physical signal, Gaussian error propagation is applied to determine the additional standard deviation of velocity and acceleration estimates resulting from the known standard deviation of the position measurement error (Gesemann 2021).

However, the statistics are influenced not only by noise but also by filtering and selection effects, such as the omission of tracks that are too short for a given filter kernel.

To avoid such selection and filter biases, it is advisable to perform numerical differentiation without prior signal filtering. Although this approach leads to a considerable overestimation of the estimated statistics, it introduces only a known offset.

As an initial example, the 5-point central difference formula is used to calculate temporal derivatives. Let s_k denote the raw particle position at time $t_k = t_0 + kh$, where h is the uniform sampling interval. The velocity at t_0 is estimated as

$$u_0 \approx \frac{s_{-2} - 8s_{-1} + 8s_1 - s_2}{12\,h}. \tag{7}$$

Assuming uncorrelated position noise with standard deviation σ_p, the variance of the velocity estimate follows from standard error propagation as

$$\sigma_u^2 = \sum_{k=-2}^{2} \left(\partial_{s_k} u_0\right)^2 \sigma_p^2 = \frac{130}{144\,h^2}\,\sigma_p^2. \tag{8}$$

This variance provides an estimate of the noise-induced uncertainty of the velocity measurement. The same error-propagation approach applies to velocities and accelerations reconstructed using TrackFit; however, in this case the resulting uncertainty depends on the weighting coefficients $\vec{c}$ used in the reconstruction.

Moreover, due to the dependence on the weights, the standard deviation of the filtered signal $\vec{s'}$ becomes time-dependent, denoted as $\sigma_{s',l}$. The corresponding propagation of uncertainty can be written as

$$\sigma_{s',l}^2 = \sum_{k=1}^{n} \left(\partial_{s_k} s'_l\right)^2 \sigma_p^2, \tag{9}$$

where s'_l denotes the filtered signal at the measurement time steps $l = 1, \ldots, n$.

Substituting the least-squares solution of Equation 4 into the sampling Equation 2 gives an expression for the filtered signal $\vec{s'}$ in terms of the unfiltered signal:

$$\vec{s'} = (-1)^i E_i \begin{pmatrix} E_0 \\ \frac{h^3}{w} E_3 \end{pmatrix}^+ \begin{pmatrix} \vec{s} \\ \vec{0} \end{pmatrix}. \tag{10}$$

The derivative of $\vec{s'}$ with respect to $s_1, \ldots, s_n$ is given by

$$\partial_{s_k} s'_l = \left[(-1)^i E_i \begin{pmatrix} E_0 \\ \frac{h^3}{w} E_3 \end{pmatrix}^+ \right]_{l,k}, \tag{11}$$

where $[\cdot]_{l,k}$ denotes the (l,k)-th entry of the corresponding matrix. Substituting this expression into Equation 9 then gives

$$\sigma^2_{s',l} = \sigma^2_p \sum_{k=1}^{n} \left(\left[E_i \begin{pmatrix} E_0 \\ \frac{h^3}{w} E_3 \end{pmatrix}^+ \right]_{l,k} \right)^2 \tag{12}$$

where the left hand side denotes the variance of the equidistantly sampled position, velocity or acceleration (depending on i) at each of the l time instants within a track. This formulation allows the temporal evolution of the measurement uncertainty to be quantified.

For sufficiently long tracks, the standard deviation of position, velocity, and acceleration is nearly constant in the central portion but slightly elevated at the track boundaries, primarily in the first and last two sampling points.

Therefore, noise effects can be effectively compensated by using the standard deviation at the midpoints of sufficiently long tracks. Accurate second order moment estimates σ^2 therefore can be computed with

$$\sigma^2 = \sigma^2_{\text{measured}} - \sigma^2_{s',l^*} \tag{13}$$

with σ^2_{s',l^*} estimated according to Equation 12 and l^* denoting half the mean track length.

When parametrizing the filter, it should be noted that filtering also removes a portion of the true signal, which can lead to an underestimation of the corrected standard deviation. Therefore, unlike in the filtering optimized for minimal Mean Absolute Error (MAE), the cutoff frequency should not be chosen at SNR = 1, but at significantly lower SNR. In this way, the additional standard deviation introduced by noise can be reduced, allowing faster convergence of the estimated statistics, while the signal loss due to filtering remains small.

In principle, the same procedure can be applied to local TrackFit, but varying cutoff frequencies across tracks result in different noise levels, preventing compensation via a single scalar excess standard deviation. In practice, using local TrackFit is not necessary here, since the additional noise contribution of global TrackFit is subtracted afterward. Consequently, the only benefit of local TrackFit would be a marginal increase in convergence speed.

With the presented approach for estimating standard deviation propagation, it can also be shown that the choice of the temporal shift τ between measurement and sampling directly affects the resulting uncertainty. The optimal shift depends on the quantity of interest: minimal standard deviation for velocity is obtained at $\tau = 0$, while the optimum for position and acceleration occurs at $\tau = 0.5$. When both velocity and acceleration are required at the same time instant, a suitable compromise must be chosen. Equation 12 can be evaluated for different values of τ to illustrate this effect.

4 Results and discussion

For clarity, results and their interpretation are presented jointly in the following section, beginning with the outcomes of the local TrackFit calibration. Since the described calibration is performed independently for each component, the subsequent analysis and the presented statistics focus on the streamwise x-component to maintain brevity.

Results of the local TrackFit calibration comprise ASDs as a function of the mean absolute velocities and accelerations (shift of the ASDs with increasing mean absolute acceleration shown in Fig. 3) as well as the resulting cutoff frequency field $f_{\mathrm{cutoff}}(\overline{|u_1|}, \overline{|a_1|})$ shown in Fig. 4.

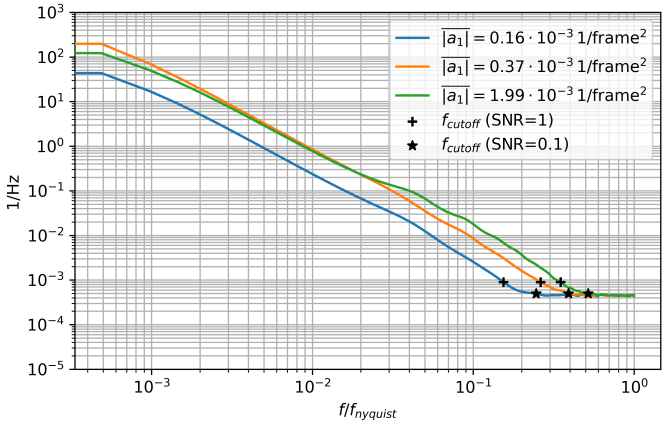


Fig. 3 ASD as a function of frequency normalized by the Nyquist frequency for bins with mean absolute velocity $\overline{|u_1|} = 0.146 \cdot 1/\mathrm{frame}$ and different mean absolute accelerations

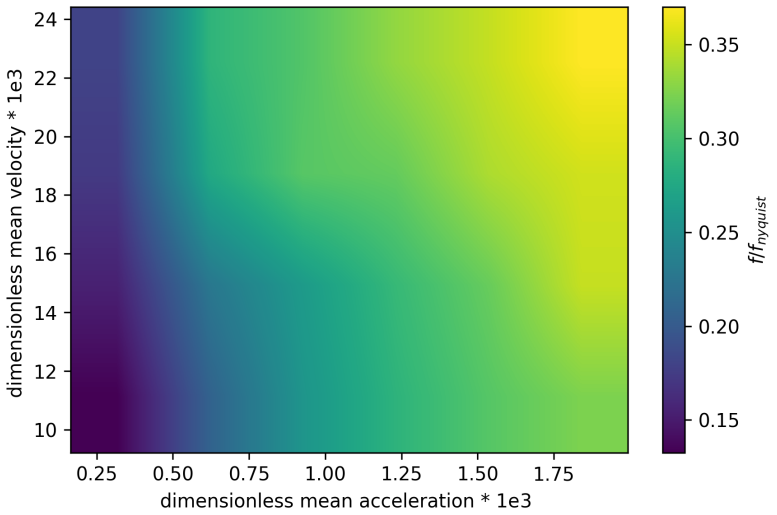


Fig. 4 Bilinear interpolation of cutoff frequencies obtained from the ASDs

The computed ASDs exhibit noticeable deviations from the idealized plateau and noise model. However, meaningful cutoff frequencies can still be derived because the automated procedure relies on the known noise level. Specifically, f_{cutoff} is defined as the frequency at which the ASD reaches twice and 1.1 times the noise level, respectively, without the need to explicitly determine the exact position of the spectral kink. Nevertheless, errors are likely to occur in bins with only a limited number of long tracks (e.g. 50 samples). Even the method proposed by (Gesemann 2021) for the robust estimation of ASDs from tracks appears to produce artifacts under such conditions. Despite these limitations, the resulting cutoff frequency field appears smooth and remains physically plausible.

In the following analyses of the accuracy of instantaneous velocity and acceleration estimates and their time-averaged statistics we examine not only the TrackFit configuration optimized for MAE performance with a low cutoff frequency (based on SNR = 1), but also a higher cutoff frequency. The latter arises when the cutoff is set directly at the knee of the ASD, corresponding to an SNR ≈ 0.1. Historically, the TrackFit cutoff frequency was determined manually at this knee point, a practice that resulted in significantly higher cutoff frequencies (in this case , f_{cutoff} increases from ~ 0.295 to ~ 0.46 for global TrackFit).

4.1 Instantaneous statistics

To assess the performance of the proposed methods for improving the accuracy of instantaneous quantities, the MAE, $\overline{e} = \overline{|X - X_{gt}|}$ between a measured quantity X and its ground truth X_{gt} is evaluated. Ground-truth velocity and acceleration are obtained from noise-free data using a 5-point central difference stencil. The results are averaged over time, spanwise and streamwise direction to obtain the MAE development

along the wall-normal direction. Results are reported separately for position, velocity, and acceleration.

First, Fig. 5 shows the position errors $\overline{e}_x$ at successive stages of the processing chain leading to locally fitted, averaged tracks. Dashed lines indicate results obtained from tracks with a tripled time-step size followed by B-spline interpolation back to the original temporal resolution. The interpolation was performed by solving Equation 5 with $f_{\text{cutoff}} = 1$ leading to a constant weight $w = \pi^3$ for all time steps within a given raw track.

Averaged tracks are obtained by averaging three interpolated tracks corresponding to the same particle. These tracks are additionally processed using local track fitting; the corresponding results are also shown in the figure. Solid lines represent raw and filtered position signals obtained without temporal subsampling.

As it is not possible to generate a cutoff frequency calibration for the averaged tracks due to the ASD shape distortion the cutoff frequency is chosen in the same way as for the tracks without averaging, rather than being based on the SNR of the averaged data. Note that the mean absolute error of the raw tracks ($\approx 3.6e^{-4}$) corresponds to the expected absolute deviation of Gaussian noise with $\sigma_x = 0.00045$.

Comparing the interpolated results with the interpolated and subsequently averaged data confirms that the expected noise reduction proportional to $1/\sqrt{3}$ is nearly achieved. Furthermore, the averaged data exhibits a comparable error to the full temporal resolution when utilizing global TrackFit. A comparison between direct global and local track fitting further shows that local track fitting yields a lower noise-induced MAE than global TrackFit in the near-wall region. Conversely, if the averaged data is used as the basis for the local track fitting, the error increases significantly away from the wall. This suggests that aliasing effects and/or the ASD shape distortion due to the interpolation degrade filtering efficiency significantly. Even in direct vicinity of the wall, where less aliasing is to be Even in the immediate vicinity of the wall, where less aliasing is expected due to the concentration of the ASD signal at lower frequencies, the performance does not exceed that achieved by track interleaving for local TrackFit reconstruction at full temporal resolution. Consequently, these data demonstrate that track interleaving combined with local TrackFit filtering is the preferred approach for high-fidelity reconstructions. Alternatively, track averaging without subsequent filtering also yields robust results, offering the distinct practical advantage of requiring no calibration of cutoff frequencies.

In the following, we compare the performance in terms of velocity and acceleration of both TrackFit approaches and both the manual way to set $f_{\text{cutoff}} \approx f_{\text{SNR}=0.1}$ at the knee point as well as utilizing $f_{\text{cutoff}} = f_{\text{SNR}=1}$.

As can be seen in Fig. 6 the local track fit yields a lower MAE in velocity and acceleration estimates compared to global TrackFit in vicinity to the wall.

This observation is consistent with the substantially lower cutoff frequencies of slowly accelerating particles, which predominantly occur in the viscous sublayer. Further away from the wall, local TrackFit configured with $f_{\text{cutoff}} = f_{\text{SNR}=1}$ only shows a slight error reduction. Hence, in homogeneous flows only a minor advantage can be expected, although both slowly and rapidly accelerating particles contribute to the overall statistics.

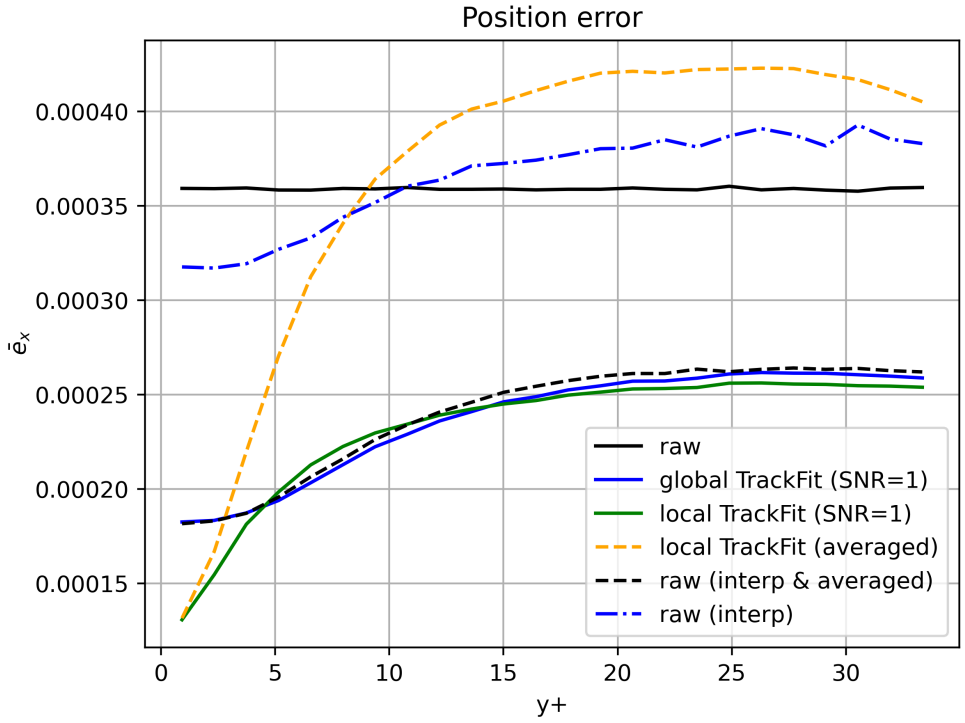


Fig. 5 Mean absolute position error $\overline{e}_x$ as a function of wall distance y^+. Dashed lines show results obtained from tracks with a tripled time step using interpolated tracks, averaged interpolated tracks, and locally fitted averaged interpolated tracks, respectively

When TrackFit is parameterized with $f_{\text{cutoff}} \approx f_{\text{SNR}=0.1}$, the same overall trend is observed; however, distinct MAE levels persist even at larger wall-normal distances. In this configuration, the local TrackFit approach yields an even greater advantage.

An open question remains as to why the global approach with SNR = 0.1 outperforms the local approach with SNR = 1 in terms of velocity error away from the wall. We hypothesize that addressing this question will also clarify why the performance advantage of local TrackFit is small in the case $f_{\text{cutoff}} = f_{\text{SNR}=1}$. We suspect that this behavior is related to the estimation of mean absolute velocity and acceleration necessary to determine the cutoff frequency for individual particles.

Specifically, the averaging window of 50 samples used to determine the cutoff frequency suppresses short-time events with high or low velocity or acceleration. As a result, locally estimated cutoff frequencies may not adequately reflect transient dynamics that are nevertheless relevant. In addition, the calibration tracks are limited to a maximum length of 50 frames, which can introduce uncertainties in the ASD estimation and thus restrict the accuracy of the cutoff frequency field $f_{\text{cutoff}}(\overline{|u_i|}, \overline{|a_i|})$ estimation.

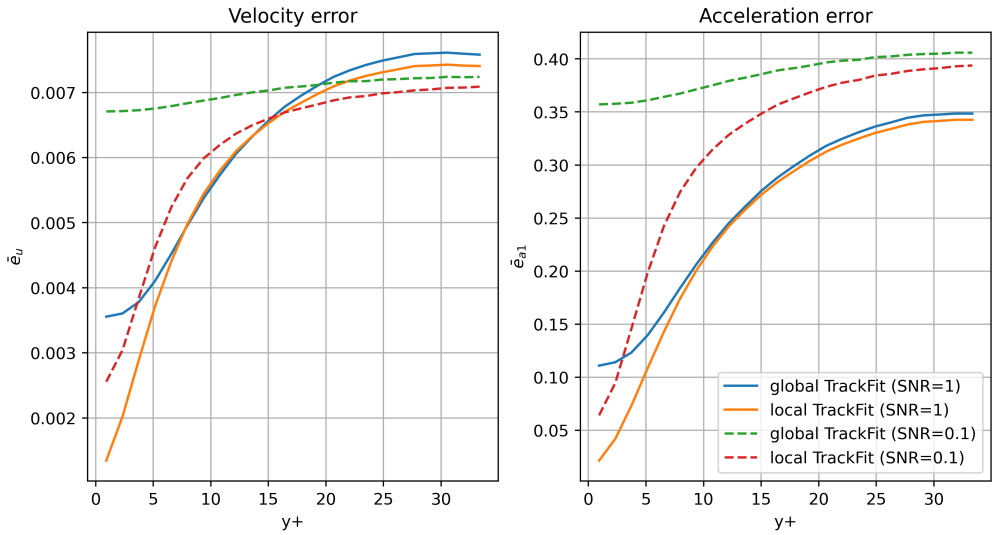


Fig. 6 Mean absolute velocity and acceleration error $\bar{e}_u$ as a function of wall distance $y+$

These limitations could be addressed by adjusting the calibration procedure. For example, instead of bilinearly interpolating the function $f_{\text{cutoff}}(\overline{|u_1|}, \overline{|a_1|})$ as in the current implementation, one could extrapolate to capture particularly high-velocity or high-acceleration events, rather than simply using the boundary value. Combined with shorter averaging windows for cutoff frequency determination during fitting, this would allow for a wider variation of cutoff frequencies and therefore a more adaptive response to the local flow conditions, potentially yielding additional accuracy improvements.

A further improvement of the local TrackFit method could be an iterative determination of both velocity and acceleration. Currently, during calibration and track fitting, the mean velocity and acceleration of each track are first estimated via smoothing with a global cutoff frequency. These estimates are then used to set the cutoff frequency for individual particles. If the initial estimates are inaccurate, the resulting cutoff frequency may be suboptimal, reducing the filter's effectiveness. To overcome this, multiple functions $f_{\text{cutoff},k}(\overline{|u_k|}, \overline{|a_k|})$ could be computed during calibration, with k as the iteration index. During track fitting, the same number of iterations would be used to update $\overline{|u_k|}$ and $\overline{|a_k|}$, potentially improving the accuracy for tracks that deviate from the global statistics.

Even without implementing this iterative approach, the relative error $\bar{e}_{rel} = \bar{e}/|X_{gt}|$ defined as the ratio between the error and the ground truth value, is already significantly reduced near the wall compared to the global TrackFit, as illustrated in Fig. 7. The relative error reduction ranges from approximately 0 to 45% for velocity and from 0 to 80% for acceleration, highlighting the substantial accuracy improvement in slow flow regions. The relative error analysis further indicates that local TrackFit with SNR = 1 appears to offer a reasonable compromise, with limited near-wall errors

and comparable errors at larger wall-normal distances relative to global TrackFit with SNR = 0.1 in the case of velocity.

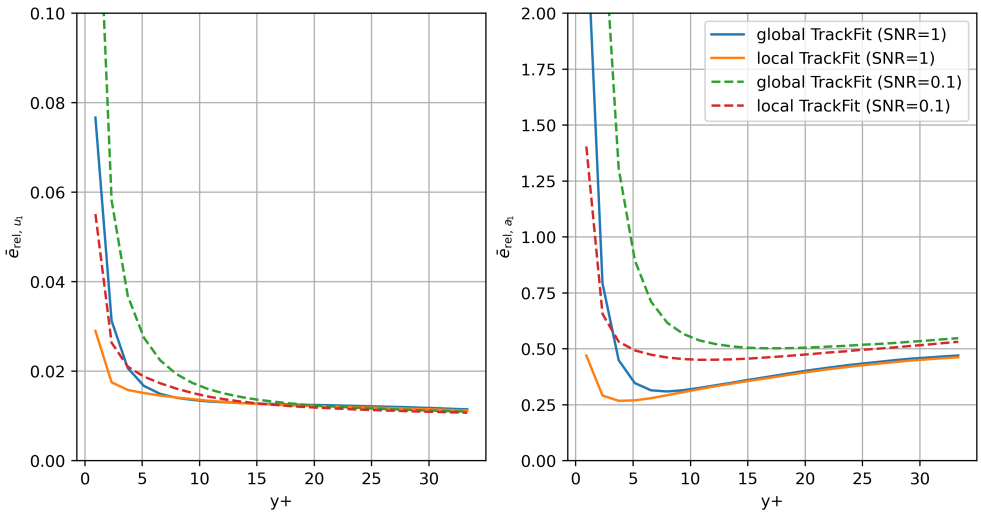


Fig. 7 Mean relative velocity and acceleration error $\bar{e}_{rel}$ as a function of wall distance $y+$

The effectiveness of the proposed local track fit approach is further supported by the results of the LPT Challenge 2025. The performance metrics reported here are taken as published by the authors on the official benchmark webpage (Leclaire et al. 2022). Fig. 8 shows RMSE results for STB using the local track fit, compared with other LPT techniques employing different polynomial-regression-based filters. Results are reported for full, half, and one-third of the original temporal resolution. The effect of high-frequency suppression becomes apparent when comparing Root mean square errors (RMSE) of position and acceleration. While the positional accuracy, both before and after fitting, is not among the best-performing methods in terms of RMSE, lower acceleration RMSE are obtained for both the full and the halved temporal resolution. This behavior can be attributed to the strong sensitivity of acceleration estimates to high-frequency noise components, which are attenuated more effectively by the local track fit than by the polynomial-regression-based filters employed in the other presented methods.

4.2 Time averaged statistics

The comparison of methods for estimating second-order moments of velocity and acceleration employs the same binning strategy as used for the instantaneous analysis. Three approaches are considered: (i) TrackFit-based estimates without noise correction; (ii) central-differences-based velocity and acceleration estimates with subsequent noise correction; and (iii) TrackFit estimates with a varying cutoff frequencies, combined with noise subtraction

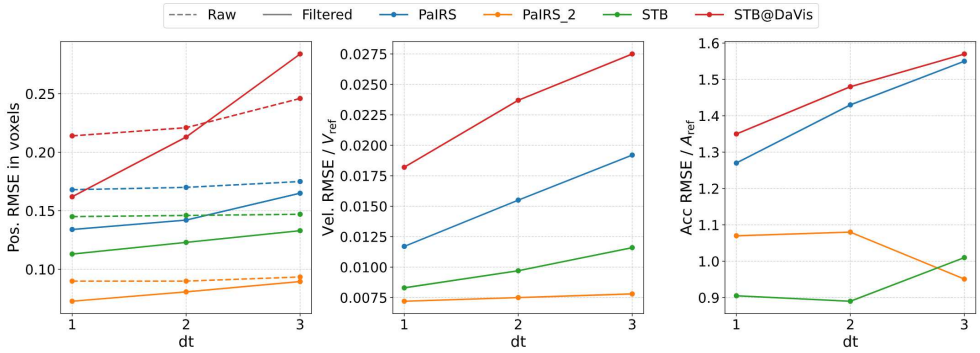


Fig. 8 RMSE of position, velocity, and acceleration, scaled by a reference velocity and acceleration, respectively. Data were obtained from LPT algorithms in a square duct flow (ONERA 2026). Local TrackFit is applied only to the "STB" case

Fig. 9 shows that, without noise correction, the velocity variance is slightly underestimated for $y^+ > 10$ when f_{cutoff} is chosen to correspond to SNR = 1. This effect is even more pronounced for acceleration, highlighting the critical impact of filtering on the accuracy of derived quantities. The observed underestimation of acceleration variance stems from the highly non-Gaussian nature of turbulent acceleration. The rare but extreme intermittent events in the heavy tails of this distribution contribute disproportionately to the total variance; for this specific dataset, the highest 5% of acceleration events account for nearly 42% of the total variance. This effect is expected to be even more pronounced in flows with higher Taylor-Reynolds numbers. To illustrate this trend, we compare our results with data from a Van Karman mixer flow at $Re_\lambda = 970$. It has been found that the acceleration PDF for this high-Re_λ case follows a stretched exponential PDF of the form

$$P(a) = C \exp\left(\frac{-a^2}{(1 + |a\beta/\sigma|^\gamma)\sigma^2}\right) \tag{14}$$

with parameters $C = 0.733$, $\beta = 0.513$, $\gamma = 1.600$, and $\sigma = 0.563$ (Mordant et al. 2004). For this specific model, the highest 5% of acceleration events accounts for as much as 63% of the total variance. Consequently, any filtering that truncates a substantial part of these high-amplitude peaks eliminates the primary physical contributors to the variance. Here, the benefit of less aggressive smoothing achieved by choosing $f_{\text{cutoff}} = f_{\text{SNR=0.1}}$ becomes apparent. While the MAE results are slightly degraded in comparison, the variances are determined more reliably for both velocity and acceleration. In the case of velocity variance estimation, the conventional TrackFit approach yields almost exact results. For acceleration, we observe an overestimation, demonstrating that the truncation bias is actually overcompensated by the residual experimental noise, which significantly inflates the variance estimation.

Because of the strong influence of extreme events, the iterative calibration method proposed in Section 4.1 has even greater potential to improve the presented statistics, as extreme values are likely particularly susceptible to distortion from suboptimal calibration.

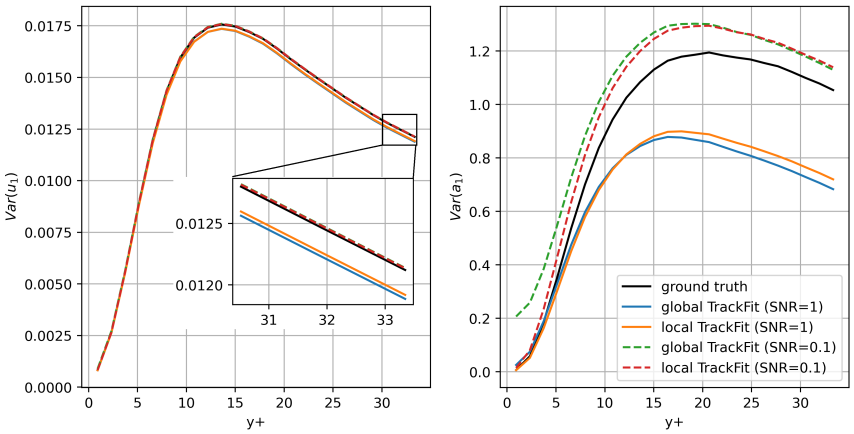


Fig. 9 Mean velocity and acceleration variance estimations using global and local TrackFit as a function of wall distance $y+$ without noise subtraction. Noise-induced acceleration variance subtracted: ≈ 4.35 (central differences), 3.72, 0.18, 0.02 (TrackFit, descending f_{cutoff}).

For robust estimation, methods incorporating noise subtraction are preferred. As shown in Fig. 10, the central differences scheme applied without prior filtering yield superior results for both velocity and acceleration moment estimation. While TrackFit combined with noise subtraction works well for velocity if a sufficiently high cut-off frequency is chosen, the acceleration plot reveals significant errors. Even when minimal smoothing is used ($f_{cutoff} = 0.8$), the variance is overestimated. This overestimation stems from the fact that the noise-induced variance increase depends on the track length. Because the mean track length varies with wall distance, incorporating local track length characteristics is essential for achieving consistently accurate results. In this study, a constant track length of 50 was assumed. This represents a realistic approximation for the near-wall region, where low flow velocities minimize the occurrence of prematurely truncated tracks caused by particles exiting the simulation domain. Crucially, in this near-wall region where the track length approximation holds true, the acceleration variance is accurately estimated, thereby validating the proposed method. Ultimately, the primary application of the described method for quantifying noise-induced standard deviation in TrackFit is not to estimate moments, but rather to improve data assimilation frameworks. Specifically, this knowledge of noise-induced uncertainty can be leveraged to optimally weight velocity and acceleration values of individual particles. While FlowFit3 currently performs well using a constant weighting between velocity and acceleration across all particles, a localized weighting strategy offers a promising pathway for further optimization. This localized weighting can be directly computed as a function of the particle's position in the track. Using central differences together with noise subtraction is the recommended approach to obtain accurate time averaged second order moment data. Nevertheless, measurement precision is inherently constrained by limited temporal oversampling relative

to the Kolmogorov timescales, causing high-acceleration events to remain unresolved. Mitigating this requires higher temporal resolution, necessitating LPT with variable timesteps to maintain tracking fidelity for particles with minimal displacement. However, such temporal refinement increases demands on high-speed imaging hardware and data storage.

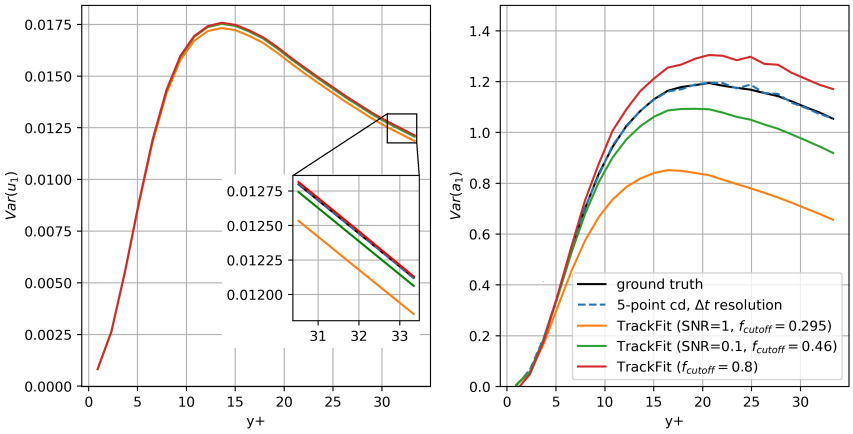


Fig. 10 Mean velocity and acceleration variance estimates obtained via central difference schemes as well as TrackFit with noise subtraction

4.3 Position prediction

The MAE of particle position predictions based on Wiener filtering, Taylor series, and quadratic extrapolation is evaluated using both the channel flow and cylinder wake datasets.In both cases, the dimensionless position signals are scaled to reflect realistic physical displacements, corresponding to pixel-based units with a mean streamwise displacement of ≈ 5 px and a maximum displacement of ≈ 10 px. Table 1 shows MAE values and the number of predictions with a deviation from the ground truth greater than 2 pixels. Beyond this threshold, the probability of successfully identifying the correct particle through shaking diminishes significantly (Schanz et al. 2022). The corresponding error histograms for the streamwise position component are shown in Fig. 11, while the histograms for the remaining two components are presented in Fig. 12.

The error histograms and MAE values obtained for both datasets demonstrate that the Taylor- and Wiener-based prediction frameworks are significantly superior to the classical quadratic extrapolation. Notably, the strong performance of the Taylor-series based prediction is remarkable given its inherent simplicity.

In the cylinder wake flow, the local calibration approach yields significantly lower MAE regardless of the chosen prediction method, whereas this trend is not observed

	Cylinder wake flow			Channel flow		
MAE in px	**x**	**y**	**z**	**x**	**y**	**z**
quadr. extrapol.	0.311	0.311	0.31	0.536	0.534	0.538
global (Wiener)	0.138	0.139	0.136	**0.298**	**0.254**	0.305
global (Taylor)	0.16	0.162	0.166	0.328	0.287	0.328
local (Wiener)	0.127	0.136	**0.098**	0.305	0.257	**0.304**
local (Taylor)	**0.121**	**0.123**	0.144	0.32	0.286	0.32
Misprediction ratio (L_2-norm $>$ 2px)						
quadr. extrapol.		**0.1**%			3.21%	
global (Wiener)		0.38%			0.21%	
global (Taylor)		0.35%			0.23%	
local (Wiener)		0.17%			**0.13**%	
local (Taylor)		0.16%			0.2%	

Table 1 MAE and misprediction ratio (proportion of predictions deviating by more than 2 pixels from ground truth) for cylinder wake flow and channel flow for quadratic extrapolation and Wiener/Taylor based prediction utilizing local and global f_{cutoff} calibration. Coordinate orientation: x (streamwise), y (wall-normal/cross-stream), and z (spanwise)

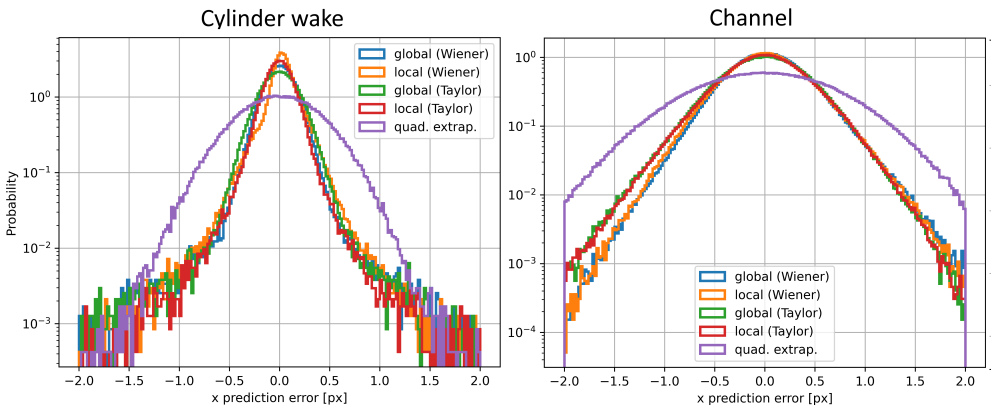


Fig. 11 Probability density functions of the prediction error for quadratic extrapolation and Wiener/Taylor based prediction utilizing local and global f_{cutoff} calibration in cylinder wake and channel flow

in the channel flow. Furthermore, the Taylor-series based method shows a slight performance advantage in the wake flow, but under performs in the channel flow configuration.

Since the investigated methods are based on fundamentally different assumptions (the Wiener framework implicitly assuming a Gaussian-distributed jerk (Kearney

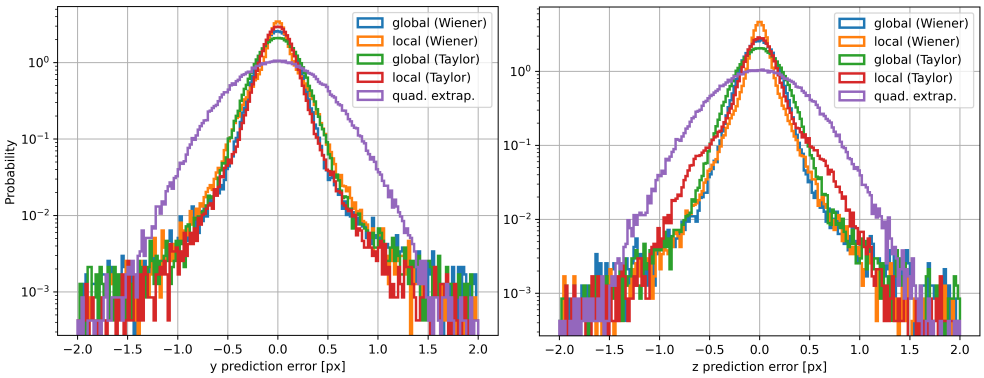


Fig. 12 Probability density functions of the prediction error of cross-stream (y) and spanwise (z) position components for quadratic extrapolation and Wiener/Taylor based prediction utilizing local and global f_{cutoff} calibration in the cylinder wake flow dataset

et al. 2026), whereas the Taylor approximation assumes constant velocity and acceleration over the time step), it is highly logical that variations in how severely these assumptions are violated across different flow regimes lead to such divergent performance. Consequently, employing a local estimation reduces the prediction error for both datasets, though the extent of this improvement varies drastically between the flow cases. Notably, the position components transverse to the downstream direction exhibit markedly different performance gains despite identical noise levels and nearly identical calibrated cutoff frequencies.

The error PDFs reveal fundamentally distinct characteristics between the quadratic extrapolation and the remaining methods. The error PDF of the quadratic extrapolation is significantly broader, yet it notably lacks the long tails observed in the other distributions, indicating a bound on extreme prediction outliers.

Furthermore, the error PDFs expose clear component-dependent discrepancies. In the channel flow dataset, both Wiener predictors exhibit a distinct tendency to overestimate the downstream position, which corresponds to a systematic bias toward higher velocities. Interestingly, this directional effect is entirely absent for the cylinder wake flow. This discrepancy underscores the necessity for a detailed investigation into the interplay between optimal predictor parametrization and the specific underlying flow properties.

Since the number of trajectories that can be successfully continued primarily depends on minimizing the fraction of particles mispredicted by more than 2 pixels (L_2-norm), this metric is also reported in Table 1.

In the cylinder wake flow, quadratic extrapolation outperformed the alternative methods, whereas it underperforms significantly in the channel flow. Potentially due to the locally lower oversampling rate relative to the Kolmogorov time scale in the channel flow, the error PDFs are inherently broader; consequently, this wider error distribution shifts the quadratic extrapolation results beyond the 2-px threshold. This suggests that, particularly at low Kolmogorov time oversampling rates, alternative prediction approaches should be pursued to prevent the premature termination of

highly accelerated tracks. Additionally, it should be noted that quadratic extrapolation simultaneously impedes the tracking process due to its large fraction of slightly mispredicted positions. Furthermore, the presented reduction of the misprediction ratio of up to 55% in the wake flow and 38% in the channel flow underscore the immense potential of local calibration in Taylor and Wiener predictors to improve overall tracking effectiveness, particularly in challenging flow regions where extreme errors would otherwise lead to premature track termination.

It is important to note that presented results were obtained with a strictly causal implementation. Despite the availability of ground truth, the predictors were fed exclusively with past data. This ensures the method's direct implementability for real experimental datasets where future information is unavailable.

The reported substantial reduction in high prediction error events can significantly enhance the performance of STB, as vortices are particularly challenging to track accurately due to their associated pressure minima, which induce highly accelerated particle motion. The elevated accelerations in these regions degrade prediction accuracy, increasing the likelihood of track termination and leading to lower resolved and biased results in these flow areas of particular interest. Consequently, subsequent data assimilation faces difficulties in reconstructing the flow, particularly in capturing the amplitude of pressure minima, due to the absence of track data in these critical regions. This issue has significant practical implications for the effectiveness of STB investigations. It directly impacts the assessment of structural loading, where pressure minima govern the forces and moments acting on bluff bodies, as well as the evaluation of lift generation in airfoils. Furthermore, improved prediction accuracy can significantly mitigate the bias in acceleration statistics caused by premature track termination, since extreme acceleration events, despite their rarity, disproportionately influence higher-order statistical moments.

A visual example of the effect of a locally adapted cutoff frequency for wiener prediction is given in Fig. 13 using two examples with high estimated absolute acceleration. For comparison, a prediction based on simple quadratic extrapolation using the three most recent timesteps as input is also presented.

For both the moderate-speed track and the slow track, significant differences in the predicted position arise, even though the local adaptation only results in a change in the cutoff frequency of 0.05 and 0.1 respectively. These examples demonstrate that small changes in cutoff frequency can make a substantial difference, even when the noise appears small compared to the position signal.

Beyond the implementation details provided in the Methods section, there is considerable additional potential for improvement: First, the synthetic datasets used here contain a limited number of tracks. In experimental measurements conducted over longer durations, a significantly larger number of tracks can be obtained, encompassing a broader range of flow conditions. This would enable more precise calibration, allowing for finer ASD binning and the inclusion of more extreme acceleration or velocity events in both Wiener prediction and required velocity/acceleration estimation.

Second, the iterative approach previously presented for calibrating local TrackFit in track smoothing, which uses acceleration and velocity estimates smoothed by local

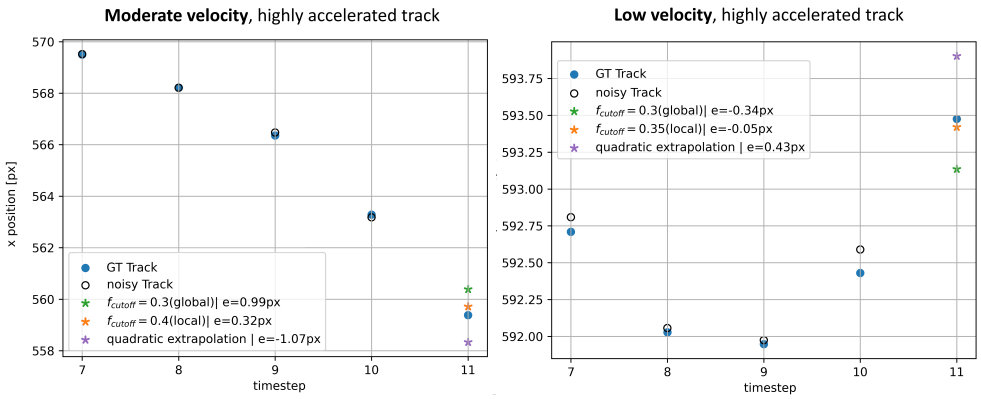


Fig. 13 Prediction of streamwise position for particles undergoing moderate and low velocity with high acceleration in the cylinder wake dataset. For each track, predictions are based on 10 input timesteps, of which only four ground truth positions are displayed as blue dots. Predictions using local and global calibration are indicated by orange and green stars, respectively. The error measure e represents the deviation between the predicted and ground truth positions

TrackFit for refined track binning, could also help identify a suitable cutoff frequency for the predictor.

Third, a positive feedback loop between the predictor's accuracy and the number of tracked high-acceleration particles could resolve the tracking problem almost entirely. This requires an iterative process: The procedure begins with an initial STB run to establish a first local calibration of TrackFit. The same dataset is then reprocessed using the improved predictor, which enhances the detection probability of highly accelerated tracks. These newly detected tracks are subsequently used to further refine the local TrackFit calibration. The process is repeated with additional STB runs, progressively reducing the number of undetected tracks. However, artificial acceleration induced by effects such as massive particle overlap may hinder accurate prediction, even when the predictor is correctly parameterized. These effects necessitate a dedicated investigation.

A combination of this approach with Variable Space-Step Shake-The-Box (VS-STB, a method to identify and correct mispredicted particles (Schanz et al. 2022)), appears particularly promising. While VS-STB can recover mispredicted particles, its performance relies on a low misprediction rate. Mispredictions result in denser residual images, which not only raise the probability of incorrect track-particle associations during VS-STB reassignment but also increase positional error. The local Wiener predictor significantly reduces the number of particles for which variable space steps must be applied, thereby facilitating a nearly acceleration-independent probability of correctly associating particles with their tracks. Ultimately, this synergy of methods could enable STB experiments with virtually no premature track terminations due to high acceleration, without the need for excessive oversampling of the Kolmogorov time scale.

5 Conclusion and outlook

This study introduces a set of complementary methods to mitigate the impact of random measurement noise in LPT data. Using synthetically generated particle tracks advected in a direct numerical simulation of turbulent near-wall channel flow, the methods were systematically validated with respect to instantaneous quantities and second order moments.

The adaptive local TrackFit approach was shown to significantly improve the estimation of instantaneous velocities and, in particular, accelerations in the near wall region by adjusting filter parameters to local flow conditions. This is beneficial in strongly inhomogeneous flows such as channel flows or jets, where conventional global filtering introduces substantial bias. A refined local TrackFit framework employing an iterative calibration strategy has the potential to further enhance performance. This superior filter performance is particularly important for advanced Lagrangian analyses like torsion statistics, since jerk is required for its computation (Hengster et al. 2024).

Beyond track post-processing, the adaptive parametrization offers clear potential for enhancing the prediction step in LPT algorithms such as Shake-The-Box, with results showing a prediction error reduction of up to 40%, with an even greater decrease in mispredictions exceeding two pixels. This improvement holds significant potential to reduce premature track terminations, thereby mitigating bias in acceleration statistics and enabling more accurate flow reconstruction, especially for pressure fields, through data assimilation.

However, a systematic investigation of Wiener filter, Taylor-series, and polynomial extrapolation-based prediction strategies, alongside their explicit dependence on flow characteristics, is required to fully realize this potential.

It has been shown that averaging of statistically independent tracks from multiple tracking passes reduces positional noise in line with theoretical expectations. Nevertheless, track interleaving with subsequent local TrackFit filtering outperforms the averaging approach, at least within the context of the channel flow test case. Both methods open up promising pathways for two distinct applications in Lagrangian particle tracking. First, they facilitate high Kolmogorov time oversampling, which is essential for accurate acceleration determination despite flow intermittency. Second, they provide a highly practical means to enhance positional accuracy for track initialization in variable-timestep frameworks, such as VT-STB (Schanz et al. 2021).

A central finding of this work is that minimizing the error of instantaneous signals demands a different filtering approach than minimizing error for statistical quantities. For the latter, we introduced a statistical noise-correction strategy that removes the contribution of measurement noise from velocity and acceleration moments using Gaussian error propagation. This approach enables noise-robust estimation of second-order statistics without track filtering and preserves the temporal resolution required to resolve small-scale dynamics.

Together, these methods demonstrate that accurate instantaneous Lagrangian velocities and accelerations require an adaptive track filtering, whereas reliable velocity and acceleration statistics can be obtained by explicitly correcting for measurement noise rather than relying on aggressive smoothing. In particular, the simplicity of the

statistical noise correction facilitates the investigation of the discrepancy between theoretical predictions of acceleration intermittency from DNS and experimental data (Buaria and Sreenivasan 2022) as well as multi-time statistics in inhomogeneous turbulent flows using comparatively simple experimental setups.

Ackknowledgements

We thank Sebastian Gesemann, who developed the TrackFit framework and pioneered its error quantification and propagation methodology, for his valuable advice on the efficient implementation of TrackFit.

Declarations

- Conflict of interest/Competing interests The authors declare no conflict of interest.
- Data availability The data will be shared upon reasonable request
- Materials availability Not applicable
- Author contribution TB wrote the main manuscript. TB, AS and DS contributed to the conceptualization of the project with the corresponding author. PG and DS contributed code. All authors read and approved the final manuscript.

References

Bosbach J, Schanz D, Godbersen P, et al (2021) Spatially and temporally resolved measurements of turbulent rayleigh-bénard convection by lagrangian particle tracking of long-lived helium-filled soap bubbles. In: 14th International Symposium on Particle Image Velocimetry, Chicago, IL

Buaria D, Sreenivasan KR (2022) Scaling of acceleration statistics in high reynolds number turbulence. Phys Rev Lett 128:234502. https://doi.org/https://doi.org/10.1103/PhysRevLett.128.234502

Buchwald T, Schanz D, Schröder A (2025) Accurate determination of acceleration statistics in a turbulent jet using locally adapted trackfit. In: 16th International Symposium on Particle Image Velocimetry, Tokyo, Japan

Cheminet A, Ostovan Y, Valori V, et al (2021) Optimization of regularized b-spline smoothing for turbulent lagrangian trajectories. Experimental Thermal and Fluid Science 127:110376. https://doi.org/https://doi.org/10.1016/j.expthermflusci.2021.110376

DiLeoni P, Agarwal K, Zaki T, et al (2023) Reconstructing turbulent velocity and pressure fields from under-resolved noisy particle tracks using physics-informed neural networks. Experiments in Fluids 64:95. https://doi.org/https://doi.org/10.1007/s00348-023-03629-4

Eilers P, Marx B (1996) Flexible smoothing with b-splines and penalties. Statistical Science 11(2):89–102

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, IL

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 Lisbon Symposium, Portugal

Godbersen P, Gesemann S, Schanz D, et al (2024) Flowfit3: Efficient data assimilation of lpt measurements. In: 21th Lisbon Symposium, Lisbon, Portugal

Hengster Y, Lellep M, Weigel J, et al (2024) Effects of anisotropy on the geometry of tracer particle trajectories in turbulent flows. European Journal of Mechanics - B/Fluids 103:284–298. https://doi.org/https://doi.org/10.1016/j.euromechflu.2023.10.008

Janke T, Michaelis D (2021) Uncertainty quantification for ptv/lpt data and adaptive track filtering. In: 14th International Symposium on Particle Image Velocimetry, Chicago, IL

Kearney G, Laurent K, Kearney R (2024) Maximum likelihood filtering for particle tracking in turbulent flows. Exp Fluids 65(24). https://doi.org/https://doi.org/10.1007/s00348-024-03765-5

Kearney GM, Laurent KM, Fardad M (2026) Maximum likelihood particle tracking in turbulent flows via sparse optimization. arXiv preprint arXiv:260222257 https://doi.org/https://doi.org/10.48550/arXiv.2602.22257

Khojasteh AR, Laizet S, Heitz D, et al (2022) Lagrangian and eulerian dataset of the wake downstream of a smooth cylinder at a reynolds number equal to 3900. Data in Brief 40:107725. https://doi.org/https://doi.org/10.1016/j.dib.2021.107725

Küchler C, Ibanez Landeta A, Moláček J, et al (2024) Lagrangian particle tracking at large reynolds numbers. Review of Scientific Instruments 95(10):105110. https://doi.org/https://doi.org/10.1063/5.0211508

Lawson J, Bodenschatz E, Lalescu C, et al (2018) Bias in particle tracking acceleration measurement. Exp Fluids 59(172)

Leclaire B, Cornic P, Champagnat F, et al (2022) A web portal for automatic performance evaluation of lagrangian particle tracking and data assimilation algorithms. In: 20th International Symposium on the Application of Laser and Imaging Techniques to Fluid Mechanics, Lisbon, Portugal

Lee M, Moser R (2015) Direct numerical simulation of turbulent channel flow up to $Re_\tau \approx 5200$. Journal of Fluid Mechanics 774:395–415. https://doi.org/https://doi.org/10.1017/jfm.2015.268

Li Y, Perlman E, Wan M, et al (2008) A public turbulence database cluster and applications to study lagrangian evolution of velocity increments in turbulence. Journal of Turbulence 9(31)

Mann J, Ott S, Anderson JS (1999) Experimental study of relative, turbulent diffusion. Tech. Rep. Risø-R-1036(EN), Risø National Laboratory, Roskilde, Denmark

Monin AS, Yaglom AM (1975) Statistical Fluid Mechanics: Mechanics of Turbulence, vol 2. MIT Press, Cambridge, MA

Mordant N, Crawford A, Bodenschatz E (2004) Experimental lagrangian acceleration probability density function measurement. Physica D: Nonlinear Phenomena 193(1):245–251. https://doi.org/https://doi.org/10.1016/j.physd.2004.01.041

ONERA (2026) 2nd lpt challenge. https://onera-flow-benchmarks.fr/en/Second_LPT_1_Results, accessed: 08 May 2026

Perlman E, Burns R, Li Y, et al (2007) Data exploration of turbulence simulations using a database cluster. Proceedings of Supercomputing SC07, ACM, IEEE

Raissi M, Perdikaris P, Karniadakis G (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. https://doi.org/https://doi.org/10.1016/j.jcp.2018.10.045

Savitzky A, Golay M (1964) Smoothing and differentiation of data by simplified least squares procedures. Analytical Chemistry 36:1627–1639

Schanz D, Gesemann S, Schröder A (2016) Shake-the-box: Lagrangian particle tracking at high particle image densities. Experiments in Fluids 57(5):70. https://doi.org/https://doi.org/10.1007/s00348-016-2157-1

Schanz D, Novara M, Schröder A (2021) Shake-the-box particle tracking with variable time-steps in flows with high velocity range (vt-stb). In: 14th International Symposium on Particle Image Velocimetry, Chicago, IL

Schanz D, Novara M, Schröder A (2022) 3d particle position determination and correction at high particle densities. In: 20th Lisbon Symposium, Portugal

Schmid M, Rath D, Diebold U (2022) Why and how savitzky–golay filters should be replaced. ACS Measurement Science Au 2(2):185–196. https://doi.org/https://doi.org/10.1021/acsmeasuresciau.1c00054

Shin H, Schröder A (2025) Robust flow field reconstruction using pinn for 3d lagrangian particle tracking. In: 16th International Symposium on Particle Image Velocimetry, Tokyo, Japan

Wiener N (1949) Extrapolation, Interpolation, and Smoothing of Stationary Time Series: With Engineering Applications. The MIT Press, https://doi.org/https://doi.org/10.7551/mitpress/2946.001.0001

Zhang J, Bhattacharya S, Vlachos PP (2020) Using uncertainty to improve pressure field reconstruction from piv/ptv flow measurements. Experiments in Fluids 61(6):131. https://doi.org/https://doi.org/10.1007/s00348-020-02974-y

Zhou K, Grauer S (2026) On the joint estimation of flow fields and particle properties from lagrangian data. URL https://arxiv.org/abs/2510.00479, arXiv:2510.00479

Zhou K, Grauer SJ (2023) Flow reconstruction and particle characterization from inertial lagrangian tracks. arXiv preprint arXiv:231109076 https://doi.org/https://doi.org/10.48550/arXiv.2311.09076