RESEARCH ARTICLE

CrossMark

# Bias in particle tracking acceleration measurement

John M. Lawson[1] · Eberhard Bodenschatz[1] · Cristian C. Lalescu[1] · Michael Wilczek[1]

Abstract

We investigate sources of systematic error (bias) in acceleration statistics derived from Lagrangian particle tracking data and demonstrate techniques to eliminate or minimise these bias errors introduced during processing. Numerical simulations of particle tracking experiments in isotropic turbulence show that the main sources of bias error arise from noise due to random position errors and selection biases introduced during numerical differentiation. We outline the use of independent measurements and filtering schemes to eliminate these biases. Moreover, we test the validity of our approach in estimating the statistical moments and probability densities of the Lagrangian acceleration. Finally, we apply these techniques to experimental particle tracking data and demonstrate their validity in practice with comparisons to available data from the literature. The general approach, which is not limited to acceleration statistics, can be applied with as few as two cameras and permits a substantial reduction in the position accuracy and sampling rate required to adequately measure the statistics of Lagrangian acceleration.

Graphical abstract Sources of bias error in Lagrangian Particle Tracking measurements are explored. Methods are presented and validated to correct acceleration statistics for the main sources of systematic errors introduced by random position error and filtering, allowing for a substantial improvement in the effective temporal resolution of particle tracking measurements.
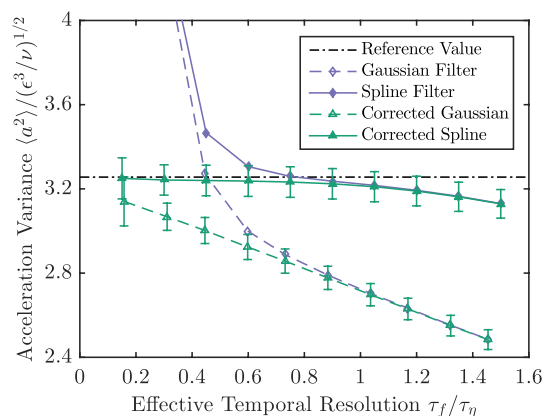
1 Introduction

The past 15 years have seen the advent of Lagrangian particle tracking (LPT) methods applied to experimental fluid mechanics. In a typical LPT experiment, time-series recordings are made of the motion of tracer particles seeded in the flow of interest. The particles are then optically tracked using standard computer vision techniques (Dracos 1996; Hoyer et al. 2005; Maas et al. 1993; Malik

✉ John M. Lawson
  john.lawson@ds.mpg.de

1   Max Planck Institute for Dynamics and Self-Organisation,
    Am Fassberg 17, 37077 Göttingen, Germany

et al. 1993; Ouellette et al. 2006). Reconstructed particle trajectories may then be numerically differentiated to obtain flow properties, such as velocity and acceleration, sampled at the particle position (Crawford 2004; La Porta et al. 2001; Mordant et al. 2004). Multiple cameras (typically three or four) can be used to extend the technique to make three-dimensional (3D) measurements (Nishino et al. 1989). LPT, therefore, readily lends itself to the examination of complex, unsteady, 3D flow phenomena from the Lagrangian frame, making it a natural choice in the investigation of problems such as turbulent transport and mixing (Holzner et al. 2008; Toschi and Bodenschatz 2009; Yeung 2002) and intermittency (La Porta et al. 2001; Voth et al. 1998, 2002). In particular, the Lagrangian acceleration plays a central role in phenomena including turbulent dispersion (Borgas and Sawford 1991; Salazar and Collins 2009) and rain initiation in clouds (Bewley et al. 2013; Falkovich et al. 2002; Shaw 2003).

The seeding concentration used in LPT is limited by the requirement to unambiguously identify and track individual tracer particles (Raffel et al. 2018). This limitation means that, instantaneously, the tracer field may not adequately sample the flow field to spatially resolve the smallest flow features present. On average, however, LPT can resolve profiles of mean flow quantities down to sub-pixel accuracy, since instantaneous flow properties may be localised to the position of individual particles (Kähler et al. 2012a). This makes LPT an ideal tool for the measurement of velocity profiles and other average field quantities, e.g., mean velocity and Reynolds stress profiles in the near wall boundary layer (Kähler et al. 2012b). In this context, systematic errors in statistics due to noise, temporal resolution, and selection biases become important, since they cannot be eliminated by taking more data.

The measurement of Lagrangian accelerations using LPT has historically been particularly challenging (Berg et al. 2009; Mann et al. 1999; Mordant et al. 2004; Voth et al. 1998, 2002). Since the position signal must be twice differentiated in time, Lagrangian acceleration measurements are very sensitive to noise. Minimising position errors beyond ~ 0.05 pixel accuracy is difficult, since they arise from a variety of sources within the measurement process, including pixelisation, image noise, quantisation errors, and overlapping particle images (Adrian and Westerweel 2011; Ouellette et al. 2006). This has led to the use of very high optical magnification to increase the position accuracy in physical space (La Porta et al. 2001; Mordant et al. 2004) and the use of cooled sensors and increased laser pulse energies to improve the image signal-to-noise ratio, which can improve position accuracy (Raffel et al. 2018). Furthermore, very high sampling rates are required to resolve very rapid fluctuations, which occur on time-scales shorter than the Kolmogorov scale (La Porta et al. 2001; Voth et al. 1998).

A common approach to mitigating noise has been to significantly oversample the position signal and subsequently apply a low-pass, finite-impulse-response filter with a large support (Crawford 2004; La Porta et al. 2001; Mordant et al. 2004; Voth et al. 2002; Xu et al. 2007). The large filter support is intended to filter out the effect of random position errors, whilst increasing the sampling rate is intended to maintain high temporal resolution. Such oversampling can lead to requirements for sampling rates on the order of tens of kilohertz in laboratory scale experiments (La Porta et al. 2001; Xu et al. 2007). Furthermore, a selection bias may be introduced when filtering is implemented as a discrete convolution (Crawford 2004; Voth et al. 1998, 2002). This requires special treatment for the ends of tracks and interpolated data and is commonly implemented by discarding data points where there is insufficient support to apply the filter (Berg et al. 2009; Mann et al. 1999). This selection bias arises because proportionally, more data are discarded for shorter tracks than longer ones, so the sample under-represents fast moving particles, which are correlated with larger accelerations (Sawford et al. 2003; Voth et al. 1998, 2002).

Once the sampling rate is fixed, the choice of filter time-scale is a compromise: short time-scales may not sufficiently filter the noise, whilst long time-scales may remove meaningful signal. Ideally, the optimal filter scale will be chosen in a range where a small change makes no difference to the statistics. This is not achieved in practice: the experiments of Berg et al. (2009), Mordant et al. (2004), and Voth et al. (1998, 2002) found no range of filter scales where the acceleration statistics are independent of scale.

As an alternative to oversampling, biases in statistical moments and probability distributions can be mitigated by making use of simultaneous, independent measurements made from two separate detectors measuring the same quantity (Crawford 2004; Mordant et al. 2004). An observation of the measurement noise can be made by taking the difference between a pair of simultaneous measurements of the same quantity. If the measurements are independent, the noise distribution may be inferred, which can be used to compensate statistical moments and probability distributions of noisy data (Stefanski and Carroll 1990).

Although numerical differentiation cannot be avoided with LPT data, the selection biases introduced whilst filtering can be mitigated. Gesemann et al. (2016) and Schanz et al. (2016) have recently popularised the use of penalised cubic B-splines for fitting noisy particle tracking data to smooth curves. The advantage of this method is that the resulting fit is twice continuously differentiable along its entire support and the curve is interpolated where data are missing. However, no systematic study has examined the influence of this approach upon acceleration statistics.

Filtering and noise are not the only sources of systematic measurement error. Other sources include preferential

concentration of tracers due to tracers not following the flow (Bec et al. 2006; Gibert et al. 2012; Monchaux et al. 2012; Toschi and Bodenschatz 2009), which introduce a selection bias effect, finite tracer size effects (Voth et al. 2002), which introduce a spatial filtering, and tracking errors (Xu 2008), which can result in spurious, large accelerations when the tracking algorithm begins following the trajectory of a different tracer. However, effective mitigations already exist for these effects: preferential concentration and finite-size effects can be avoided by using smaller, density-matched tracers which more faithfully follow the flow and tracking errors can be mitigated by reducing the seeding density or increasing the measurement frequency.

In this paper, we present a suite of methods to correct Lagrangian measurements for biases introduced by noise and filtering effects. These are based on the use of smoothing splines and independent measurements to correct noise biases. The methods are tested using numerical simulations of LPT measurements and via application to real experimental data. This allows us to realise a substantial reduction in the temporal resolution and position accuracy required to adequately measure the statistics of Lagrangian accelerations.

The paper is organised as follows. In Sect. 2, we outline a suite of techniques to correct LPT data for systematic measurement errors introduced by noise and filtering effects. Subsequently, Sect. 3 outlines the generation of the experimental, numerical simulation, and synthetic particle tracking data sets used to test these methods. The bias correction methods are validated with synthetic particle tracking data in Sect. 4. We then apply and further validate these techniques with real experimental data in Sect. 5. Concluding remarks are provided in Sect. 6.

2 Methodology

We now present a suite of methods which may be used to compensate or minimise systematic biases introduced by noise and filtering effects. We first describe how simultaneous measurements can be used to correct statistics of noisy data in Sect. 2.1 and outline how such simultaneous measurements can be applied to LPT data in Sect. 2.2. We then recap the use of filtering methods in Sect. 2.3 and identify the penalised smoothing spline as a means to reduce selection bias when filtering data.

2.1 Noise

Simultaneous measurements are a powerful tool to eliminate biases introduced by noise. Suppose that one has two simultaneous, independent measurements $a_1 = a + \gamma_1$ and $a_2 = a + \gamma_2$ of a quantity a (say, acceleration) with random noise errors γ_1 and γ_2. We can then define sums and differences of these quantities:

$$
\begin{aligned}
\overline{a} &= \frac{1}{2}(a_2 + a_1) = a + \overline{\gamma} \\
\overline{\gamma} &= \frac{1}{2}(\gamma_2 + \gamma_1) \\
\widehat{a} &= \frac{1}{2}(a_2 - a_1) = \frac{1}{2}(\gamma_2 - \gamma_1) = \widehat{\gamma},
\end{aligned}
\tag{1}
$$

which are denoted with an overbar ($\overline{a}$) or hat ($\widehat{a}$), respectively. From this, the pth moment of a is given by the following:

$$
\langle a^p \rangle = \langle \overline{a}^p \rangle - \sum_{k=1}^{p/2} \binom{p}{2k} \langle a^{p-2k} \rangle \langle \widehat{a}^{2k} \rangle,
\tag{2}
$$

where we have assumed the independence of a, γ_1, and γ_2 and that γ_i are symmetrically distributed. The key ingredients are that $\overline{\gamma}$ and $\widehat{\gamma}$ have the same distribution, are independent of a, and their odd moments vanish. We shall review the appropriateness of these assumptions in Sect. 2.2. The first three even moments of a are given by

$$
\begin{aligned}
\langle a^2 \rangle &= \langle \overline{a}^2 \rangle - \langle \widehat{a}^2 \rangle \\
\langle a^4 \rangle &= \langle \overline{a}^4 \rangle - 6\langle a^2 \rangle \langle \widehat{a}^2 \rangle - \langle \widehat{a}^4 \rangle \\
\langle a^6 \rangle &= \langle \overline{a}^6 \rangle - 15\langle a^2 \rangle \langle \widehat{a}^4 \rangle - 15\langle a^4 \rangle \langle \widehat{a}^2 \rangle - \langle \widehat{a}^6 \rangle.
\end{aligned}
\tag{3}
$$

The insight here is that the measured distribution of $\widehat{a}$ constitutes an observation of the distribution of $\overline{\gamma}$, which can be used to compensate systematic errors in statistics of $\overline{a}$. Despite the simplicity of this approach, it is rarely used within the particle tracking literature. We are only aware of the experiments of Voth et al. (2002) and later Crawford (2004) and Mordant et al. (2004), who used a unique setup of 1D silicon strip detector sensors to make pairs of independent, 2D position measurements, which were used to estimate the error in their acceleration variance measurements.

The magnitude of the noise can be described by the signal-to-noise ratio, defined as

$$
\text{SNR} = -10 \log_{10}(\langle a_m^2 \rangle / \langle a^2 \rangle - 1),
\tag{4}
$$

and is measured in decibels (dB). Here, $\langle a_m^2 \rangle$ represents some measured variance and we have assumed the measurement noise is independent. Ordinarily, one cannot directly measure this quantity, because the ground-truth variance $\langle a^2 \rangle$ is not known. However, since (3) may be used to estimate the ground truth $\langle a^2 \rangle$, an estimate of the signal-to-noise ratio can be made.

The noise-correction procedure can be extended to the level of the probability density function (PDF). Let us write the PDF of a random variable a as f_a. Since $\overline{a}$ is the sum

of two measurements of a with independent noise, we can alternatively write $\bar{a} = a + \bar{\gamma}$. Since a and $\bar{\gamma}$ are also independent, the PDF of $\bar{a}$ is given as $f_{\bar{a}} = f_a \star f_{\bar{\gamma}}$, i.e., the convolution of f_a with $f_{\bar{\gamma}} = f_{\hat{\gamma}} = f_{\hat{a}}$. Then, we can make use of deconvoluting kernel density estimators (Stefanski and Carroll 1990; Wang and Wang 2011) to deconvolve $f_{\bar{a}}$ with $f_{\hat{a}}$. This technique is well known within the statistics literature, but has surprisingly not been adopted by the particle tracking community, despite its obvious suitability.

The deconvolution method is more readily understood in terms of characteristic functions of random variables. The characteristic function of variable a with sample space frequency t is defined as $\phi_a(t) \equiv \langle e^{ita} \rangle = \int da f_a(a) e^{ita}$ and is the Fourier transform of f_a. Under the preceding statistical independence and symmetry assumptions, we have:

$$\phi_{\bar{a}}(t) \equiv \langle e^{it(a+\bar{\gamma})} \rangle = \langle e^{ita} \rangle \langle e^{it\bar{\gamma}} \rangle = \phi_a(t)\phi_{\hat{a}}(t). \tag{5}$$

This allows us to correct our measured distribution $f_{\bar{a}}$ using the distribution of the error difference $f_{\hat{a}}$. Naïvely, we might calculate f_a from $\phi_a(t) = \phi_{\bar{a}}/\phi_{\hat{a}}$. However, the factor by which the characteristic function $\phi_{\bar{a}}$ must be corrected by becomes large for large sample space frequencies t, so we introduce a kernel function $\phi_K(t)$, e.g., (Wang and Wang 2011):

$$\phi_K(t) = \begin{cases} (1-t^2)^3 & -1 \leq t \leq 1 \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

with a characteristic bandwidth h to define our deconvolution kernel $\phi_L(t)$:

$$\phi_L(t) = \frac{\phi_K(t/h)}{\phi_{\hat{a}}(t)}. \tag{7}$$

The corrected PDF $\widetilde{f}_a$ can then be obtained from the inverse Fourier transform of $\widetilde{\phi}_a(t) = \phi_{\bar{a}}(t)\phi_L(t)$. The scale of the deconvolution kernel is given by the parameter h. The choice of kernel scale comes down to a compromise between bandwidth and statistical uncertainty. Wang and Wang (2011) provide a number of methods to choose the bandwidth; we also discuss the choice of bandwidth in Sect. 4.1.

In practice, we obtain $\phi_{\bar{a}}$ and $\phi_{\hat{a}}$ from the discrete Fourier transform of fine-grained histograms of $\bar{a}$ and $\hat{a}$. Subsequently, we obtain the corrected PDF $\widetilde{f}_a$ from the inverse transform of $\widetilde{\phi}_a$. Standardised PDFs are always scaled by $\sigma_a \equiv \langle a^2 \rangle^{1/2}$ with $\langle a^2 \rangle$ obtained from (3).

2.2 Making independent measurements

To use the noise-correction techniques described in Sect. 2.1, we need to make simultaneous independent measurements of velocity and acceleration. In this section, we outline a novel approach to doing so, by making partial measurements of the velocity or acceleration from independent cameras.

The image space velocity $\dot{y}_i \in \mathbb{R}^2$ of a particle in camera i is given by

$$\dot{y}_i = J_i(x)u, \tag{8}$$

where $u \in \mathbb{R}^3$ is the particle's velocity and $J_i(x) = \partial y_i / \partial x$ is the Jacobian of its projected position $y_i \in \mathbb{R}^2$ in camera i with respect to its position $x \in \mathbb{R}^3$. Treating the projection as locally linear, one can also approximate the image space acceleration as $\ddot{y}_i = J_i(x)a$, where $a \in \mathbb{R}^3$ is the particle acceleration. Since these expressions have the same form, we will continue our analysis just for the velocity.

The Jacobian has a null space vector $n_i = n_i(x)$, which means that the component of u parallel to n_i cannot be determined from (8) alone. For a pinhole camera model, this is the component parallel to the viewing direction. We can, however, uniquely solve (8) for $\tilde{u}_i = u - (u \cdot n_i)n_i$ with the constraint $\tilde{u}_i \cdot n_i = 0$ to obtain a projection of the velocity from that measured in a single image. This allows the components of u orthogonal to n_i to be measured from a single camera. With this formulation, we see that a pair of cameras i, j can make independent measurements of the same velocity component in the direction $n_i \times n_j$, which is orthogonal to n_i and n_j.

We can combine measurements from selected sets of cameras to measure all three components by solving (8) in the least-squares sense,[1] as given by:

$$u = \left(\sum_i J_i^T J_i \right)^{-1} \sum_i J_i^T \dot{y}_i. \tag{9}$$

In this way, we can make independent measurements of the velocity from different sets of cameras. For example, with three cameras, we could use one principal camera and a second pair to get independent measurements of two components simultaneously. With four cameras, simultaneous independent measurements of all three components can be made using two pairs.

We note, in passing, that it may also be desirable to measure velocity increments $\Delta u = u(x + r, t + \tau) - u(x, t)$. With two or more cameras, this is no problem. With a single camera, technically only the component parallel to $n_i(x) \times n_i(x + r)$ may be measured. For the pinhole camera model, this is perpendicular to the separation vector r, i.e., it corresponds to a transverse velocity increment. However, if the measurement volume is far from the camera pinhole, $n_i(x)$ only varies slightly within the volume. In this case, two components of Δu can be measured, which are approximately perpendicular to $n_i(x + r/2)$.

---

[1] The problem is posed as finding an approximate solution for the overdetermined linear system $Au = b$ by minimising $\|Au - b\|_2$, where $A^T = [J_1^T J_2^T \ldots]$ and $b^T = [\dot{y}_1^T \dot{y}_2^T \ldots]$.

The image space velocity $\dot{\mathbf{y}}_i$ can be obtained from numerical differentiation of the image space position. Where necessary, it should be obtained from independent interpolations in image space to avoid introducing correlations in errors across cameras. For particle tracking codes which do not provide independent measurements of particle image position, image space coordinates of particles might be obtained separately and associated with tracks in a post-processing step.

The Jacobian can be readily computed from the camera model. For the purpose of illustration, we demonstrate this for the pinhole camera model used presently. In this case, the transformation mapping x to y_i and its Jacobian $\mathbf{J}_i(x)$ can be written in the form (Hartley and Zisserman 2003):

$$y_i = \frac{\mathbf{T}_i(x - x_{c,i})}{\mathbf{\Lambda}_i(x - x_{c,i})} \quad \text{and} \quad \mathbf{J}_i = \frac{\mathbf{T}_i - y_i \mathbf{\Lambda}_i}{\mathbf{\Lambda}_i(x - x_{c,i})}. \tag{10}$$

Here, $x_{c,i}$ is the location of the camera pinhole in object space, whilst $\mathbf{T}_i$ and $\mathbf{\Lambda}_i$ are 2×3 and 1×3 matrices parameterising the camera's orientation, magnification, and distortion.

Our main assumption is that measurements of image space velocity and acceleration in separate cameras constitute independent measurements with independent errors. In a well-calibrated LPT system, these velocity and acceleration errors arise from errors in position measurement. The independence assumption is reasonable for random errors in particle position, such as those introduced by image noise, overlapping particle images, irregularities in particle images, and multiple scattering (Adrian and Westerweel 2011). This is because these error sources occur separately for each camera and independently of particle position. On the other hand, systematic errors in the measurement of particle position such as peak-locking may violate the independence criterion, since this effectively introduces a quantisation error into the velocity or acceleration. Thus, the method is restricted to the case where random errors in particle position dominate bias errors, which can be achieved by ensuring that the particle image diameter is sufficiently large (Prasad et al. 1992).

The independence assumption is also violated when using the Shake The Box technique (Schanz et al. 2016), since particle positions are jointly optimised across cameras. This might be mitigated using independent sets of cameras to perform separate optimisations of particle positions during the "shake" step or making entirely separate measurements which are combined later. Independence may also break down when particles become very close to one another, since overlapping effects may become correlated across cameras. A further potential source of correlated error may present when tracer particles have a non-uniform shape or illumination, where apparent motion of the glare point relative to the particle center may be interpreted as particle motion. We neglect the error in $\mathbf{J}(x)$ that contributes to error in u.

2.3 Track filtering and interpolation

The conventional approach to obtaining the velocity and acceleration of fluid tracers from particle tracks is to apply a finite-difference method in combination with some level of smoothing filter (e.g., Crawford 2004; Mann et al. 1999; Mordant et al. 2004; Voth et al. 2002[2]). In our analysis, we pick a Gaussian-weighted least-squares approximation to the position, velocity, and acceleration by convolving with a set of discrete filter kernels. For some filter of scale w and support of l samples, the kernels

$$h_n = C_1 \exp(-n^2/w^2) \tag{11}$$

$$h'_n = n C_2 \exp(-n^2/w^2) \tag{12}$$

and

$$h''_n = C_3 (2n^2/w^2 - 1) \exp(-n^2/w^2) - C_4 \tag{13}$$

can be convolved with the discretely sampled position signal $\mathbf{x}(t_n)$ to obtain the filtered position, velocity, and acceleration, respectively. We choose a support of $l = 3w$ with $n \in -l/2, \ldots, l/2$. The coefficients $C_1 \cdots C_4$ are chosen to satisfy the normalisation conditions $h_n \star 1 = 1$, $h'_n \star 1 = 0$, $h''_n \star 1 = 0$, $h'_n \star (n\Delta t) = 1$ and $h''_n \star (n\Delta t)^2 = 2$, where Δt is the sampling period. More details on the choice of this filter can be found in Mordant et al. (2004).

A limitation of this approach is that the filtered quantities are undefined near where data are missing (e.g., where tracks have been reconnected) and the ends of tracks, where samples are not available to apply the convolution kernel in (11). As we show in Sect. 4.1, simply ignoring these data leads to a selection bias in the acceleration statistics.

An alternative approach to filtering noisy data is to fit a smoothing spline (Eilers and Marx 1996; Gesemann et al. 2016). The idea is to fit a spline curve $g(t)$ to the data (t_i, y_i), $i = 1 \ldots m$ which makes a trade-off between the closeness of the fit and the roughness of the curve. Gesemann et al. (2016) proposed the use of a cubic spline fit with a third-order roughness penalty, which minimises the following objective function:

$$\sum_{i=1}^{m} (g(t_i) - y_i)^2 + f_s \lambda \int_{t_1}^{t_m} g^{(3)}(t)^2 \, \mathrm{d}t. \tag{14}$$

Here, $g^{(3)}(t)$ is the third derivative of the curve, λ parameterises the level of smoothing, and $f_s = 1/\Delta t$ is the sample rate. Numerically, we implement this by performing a penalised

---

[2] Piecewise polynomial fitting, as used by Voth et al. (2002), effectively implements finite differencing and smoothing in one step.

Fig. 1 **a** Illustration of the Lagrangian Exploration Module and **b** a cross section of the experimental measurement volume. The magenta and blue isosurfaces of constant seeding density encapsulate 95% and 99% of the measured data, respectively. The acceleration of tracers is only sampled inside the green ellipsoidal volume. Pairs of independent measurements of the acceleration component are made in the direction shown by the black arrow
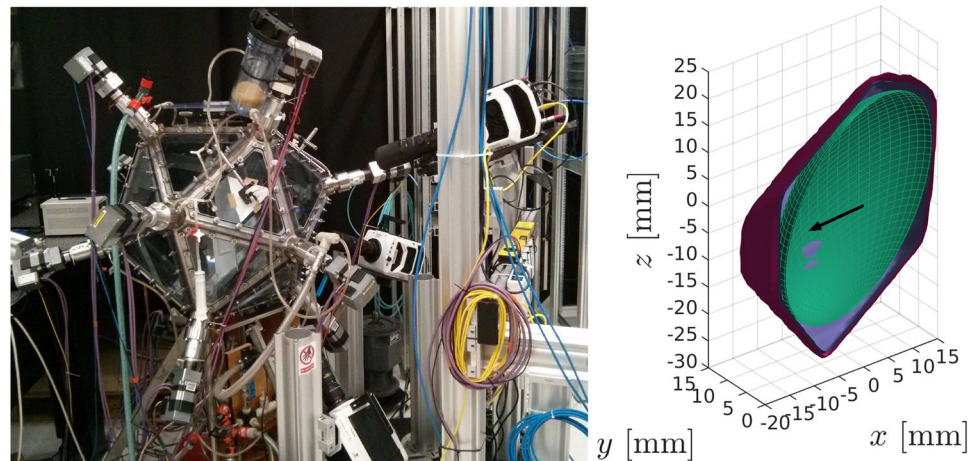
(a) Experimental Setup **(b)** Measurement Volume

linear fit of a set of $m + 2$ B-spline curves to the data. Details of how to do this can be found in Eilers and Marx (1996).

This approach has the advantage of generating a smooth representation of the underlying data which is continuous in its second derivative over the entire length of the data. As such, it may interpolate missing data in particle tracks, e.g., where tracks have been reconnected. Whilst end effects are present (the smoothing criterion results in $g^{(3)} = 0$ at the ends), this effect is much less severe than for a simple convolution filter, since the data at the ends are still represented.

The filter parameter λ determines the trade-off between the smoothness and the closeness of the fit. When $\lambda = 0$, the chosen spline interpolates the data exactly. As $\lambda \to \infty$, the fit becomes a linear least-squares regression to a quadratic polynomial. The frequency response approximates a sixth-order low-pass Butterworth filter with cut-off frequency $f_c = 1/(2\pi\lambda^{1/6})$.

To make a fair comparison between filters, we compare their performance in terms of their Equivalent Noise Bandwidth (ENBW). The ENBW is the bandwidth of the equivalent brick-wall filter with the same integrated (white) noise power and may be calculated from the impulse response h_n as follows (Elliott 1987):

$$f_{\text{ENBW}} = \frac{f_s}{2} \frac{\sum_n h_n^2}{(\sum_n h_n)^2}. \tag{15}$$

The ENBW of the Gaussian and spline filters are determined by the parameters (l, w) and λ, respectively.

3 Experimental and numerical data sets

To test the methods described in Sect. 2, we have conducted laboratory particle tracking experiments and numerical simulations of particle tracking. Our numerical simulations

consist of two aspects: direct numerical simulation (DNS) of tracer particles in homogeneous isotropic turbulence and the subsequent simulation of experimental measurement of particle tracks. We first describe our experimental measurement, and then the DNS and the synthetic particle tracking.

3.1 Laboratory experiment

We conducted Lagrangian particle tracking experiments of low Stokes number tracer particles in deionised water in the homogeneous, isotropic turbulence generated in our Lagrangian Exploration Module (LEM) facility, illustrated in Fig. 1a. We refer the reader to Zimmermann et al. (2010) for a full description of the facility and provide only a brief outline here.

The LEM consists of an icosahedral tank with transparent polyacrylamide windows and impellers at each of its 12 vertices. This configuration allows us to generate a turbulent flow at the center of the tank which is statistically homogeneous and isotropic, with a mean flow speed below 10% of the fluctuating velocity. The turbulence intensity and hence Reynolds number are adjusted by varying the rotation rate of the impellers. In the present experiments, we chose an isotropic forcing with all impellers rotating at the same frequency (between 60 and 960 rpm). The temperature of the deionised water is maintained at 20 °C by cooling plates at the top and bottom of the tank, whose cooling power is regulated by a closed-loop feedback controller. This maintains the kinematic viscosity of the water at $\nu = 1.004 \, \text{mm}^2/\text{s}$ and mass density at $\rho_f = 0.997 \, \text{kg/L}$.

Three high speed cameras (Phantom V2511, Vision Research) equipped with Nikon 200mm macro lenses and $2 \times$ teleconverters observe a measurement volume at the center of the LEM. The camera configuration and measurement volume are illustrated in Fig. 1. The region

Table 1 Characterisation of experimental data sets

R_λ	f_I (Hz)	u' (mm/s)	ϵ (cm²/s³)	L_{int} (mm)	τ_η (ms)	St	$f_s\tau_\eta$	η (px)
109	1	14.6	0.568	55	133	0.001	92.3	11
203	3	47.7	18.8	58	23.1	0.006	64.2	4.6
352	8	135	402	61	5.00	0.013	62.5	2.1
438	12	207	1420	62	2.66	0.052	66.4	1.6
504	16	276	3390	62	1.72	0.080	43.0	1.3
190	–	46.9	19.9	52	22.4	0	98.1	4.5

Details of the synthetic particle tracking simulation, detailed in Sect. 3.3, are also included in the last row

mutually visible across all cameras spans approximately $37 \times 26 \times 42$ mm. A 70W self-built pulsed Nd:YAG laser is coupled to beam forming optics to illuminate polystyrene tracer particles with diameter $d_p = 40\,\mu$m and mass density $\rho_p = 1.05$ kg/L (TS40, Microbeads AS). For each experimental condition in Table 1, we acquired 4000 independent time-series with O(1000) particles per image, corresponding to a seeding concentration around 47 particles per cm³. Each 3500 frame movie is downloaded over 10GBit Ethernet and saved in a sparse format by retaining only pixels and their neighbours with brightness above a specified threshold. Sparsification reduces the storage requirement by 90–95% over uncompressed images. The fast download and sparsification process has allowed us to acquire very large data sets with O(10^{10}) datapoints per set.

We present data collected at five different Reynolds numbers, detailed in Table 1. The integral length-scale $L_{int} = u'^3/\epsilon \simeq 60$ mm, defined in terms of the mean dissipation rate ϵ and the root-mean-square velocity fluctuation u', is approximately independent of the impeller rotation rate f_I. The mean dissipation rate was obtained from measurements of the velocity–acceleration structure function (Mann et al. 1999). The sampling rate f_s is chosen to be 43–92 times faster than the Kolmogorov frequency, which is necessary to capture the heavy tails of the acceleration distribution (Crawford 2004; Mordant et al. 2004). This results in a relatively small RMS inter-frame particle displacement of around 0.5 pixels. The Stokes number St $\equiv \tau_p/\tau_\eta$ is a measure of the response time $\tau_p = d_p^2/12\beta\nu$ of our tracers in comparison to the Kolmogorov scale $\tau_\eta = (\nu/\epsilon)^{1/2}$, where $\beta = 3\rho_f/(2\rho_p + \rho_f)$. Since the Stokes numbers achieved are small, we expect filtering effects due to particle size to be negligible (Lalescu and Wilczek 2018a; Voth et al. 2002). However, we note that preferential concentration effects may influence the acceleration statistics at the highest Reynolds numbers (Bec et al. 2006; Gibert et al. 2012).

Particle tracking is performed using an in-house code which implements a version of the predictor–corrector tracking algorithm described in Ouellette et al. (2006). We describe the procedure here only briefly. In each frame, particle images are identified using a local maximum criterion and their image centers obtained using a three-point Gaussian fit. Then, the existing particle tracks are extrapolated using a quadratic polynomial fit to six points at the end of the track, to predict the new image space position of the particle in each camera. Tracks are then associated with particle images if possible and a triangulation is performed to recover the new object space position. Tracks are terminated if a suitably accurate match (≤ 1 px triangulation error) is not found or matching is ambiguous. Due to the small interframe displacement, the tracking process is very robust: the largest tolerated prediction error is approximately half a particle diameter. New tracks are initiated by stereo-matching and triangulating unmatched particles and are extrapolated with zero displacement, but larger prediction tolerance. Tracks with fewer than three points are extrapolated with a linear fit. The procedure is iterated until all frames have been processed. Since tracks are frequently interrupted, we reconnect track segments based on their proximity in position-velocity space using the method described in Xu (2008). After reconnection, the average track length ranges between 302 and 493 frames, depending on the data set.

Due to the finite size of the measurement volume, our measurement of the distribution of particles is sharply truncated near the edges of the measurement volume. This is illustrated in Fig. 1b, which shows isosurfaces of the average seeding density within the measurement volume. To avoid selection biases associated with particles entering and leaving near the edges of the measurement volume, we only sample the statistics of tracers when they are inside the ellipsoidal region shown. The principal diameters of this ellipsoid measure $20.0 \times 27.0 \times 40.8$ mm³.

3.2 Direct numerical simulation

To provide data for our simulation of our laboratory particle tracking experiment, we ran a direct numerical simulation of tracer particles in forced, homogeneous, isotropic turbulence at $R_\lambda = 190$ in a 1344^3 periodic cubic box of side length 2π. A standard pseudo-spectral scheme was used to solve the Navier–Stokes equations in their vorticity formulation with statistical stationarity maintained by means of a large-scale
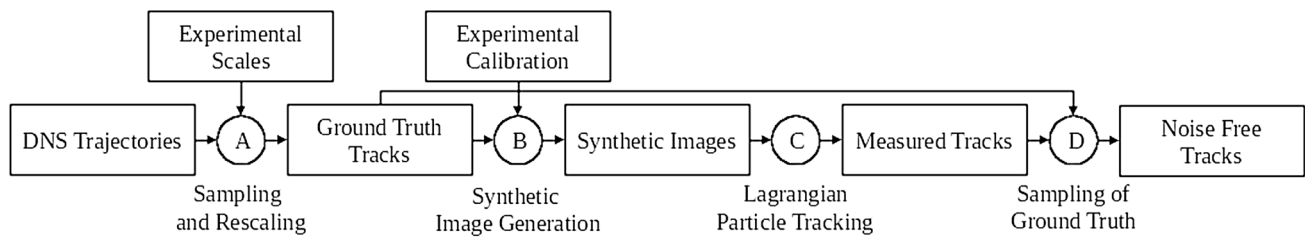
Fig. 2 Data flow of synthetic particle tracking procedure

(a) 3D Projection

(b) 2D Synthetic Image
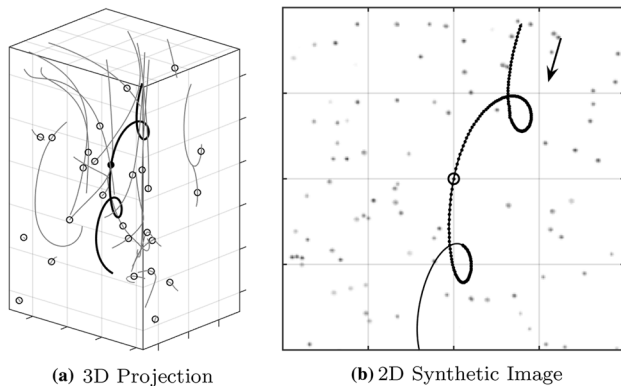
Fig. 3 Snapshot of a simulated particle, its neighbours, and their trajectories. **a** 3D projection of ground-truth particle track (bold, filled) and its neighbours (open, grey). The seeding density and Kolmogorov scales are comparable to experiment. **b** Corresponding synthetic particle image for camera 3. The large circle indicates the position at the reference time, whilst small dots indicate the image position recovered from LPT at every fourth timestep. The measured track is incomplete, because the particle leaves the measurement volume. Grid markings in both figures correspond to 64 pixels (2.1 mm)

band-passed Lundgren forcing in the wavenumber range [1.5,3]. This resulted in an integral length-scale (based on the velocity autocorrelation) over 11 times smaller than the box size, which helps to minimise effects of the periodic boundary conditions on flow statistics. The chosen spatial resolution ($k_{max}\eta = 2$, where k_{max} is the maximum resolvable wavenumber) ensures that the small scales are adequately resolved to capture extreme events.

After achieving statistical stationarity, 10^7 tracer particles were introduced and advected with the flow and their position, velocity, and acceleration were recorded over 0.5 integral time-scales ($10.4\tau_\eta$). We stored the tracer state at $\sim 0.01\tau_\eta$ intervals in order to recover the extreme acceleration events from the position signal. Additional care was taken over the integration of the tracers: cubic splines were used to interpolate the underlying velocity fields, and time-stepping was performed using a fourth-order Adams–Bashforth method. This effort has ensured that the Lagrangian velocity and acceleration statistics, as obtained from finite differences of trajectories, are in

good agreement with those sampled from the fields. Further details of the solver can be found in Lalescu and Wilczek (2018b).

3.3 Synthetic particle tracking experiment

Trajectories from the DNS were used to simulate the experimental acquisition of particle tracks. The data flow is outlined in Fig. 2. There are four steps, which we now outline in detail.

The first step (A) is to sample subsets of the DNS trajectories (corresponding to different sub-volumes from the full simulation) and rescale these from code units to physical units to match the Kolmogorov scales and seeding density of the experiment (see Table 1). This defines the "ground truth" of the synthetic measurement, a sample of which is shown in Fig. 3a. The second step (B) is to generate particle tracking images based on these ground-truth data. For this, we projected the positions of tracers in the experimental geometry onto images using the same geometric camera calibration used in the experiment. A procedure similar to the reprojection step of the Shake The Box algorithm was used to render particle images (Schanz et al. 2016). These are simulated using three-parameter Gaussian intensity profile, which was obtained from the experimental calibration of the optical transfer function, following Schanz et al. (2013). Images are then quantised to 8-bit resolution and saved using the same sparse compression format as the experimental data. The procedure effectively simulates an experimental measurement of tracer particles from DNS. A sample synthetic particle image is shown in Fig. 3b. In total, we generated 5.17×10^6 such images, corresponding to 5040 subsets of tracers over 1025 timesteps.

Using a procedure similar to (Voth et al. 2002), we then processed these synthetic particle images with the same particle tracking and post-processing toolchain as experimental data described in Sect. 3.1 (step C). This produces an experimental sampling of the ground-truth data set, as illustrated in Fig. 3b. The synthetic data set is matched to the $R_\lambda = 203$ experiment in terms of Reynolds number, scale, measurement geometry, optical properties, and processing scheme. Amongst others, our modeling simplifications overlook

effects such as particle inertia, collisions, polydisperse particle sizing, irregular particle shape, illumination instability, shadowing, image noise, and multiple scattering. These effects are expected to impair measurement quality. As such, the synthetic data sets should be considered as a "best case scenario" measurement analogous to our experimental cases, which incorporate the major sources of bias error. Crucially, the synthetic data set captures the bias effects which we identify in Sect. 1. The availability of ground-truth information permits a quantitative test of the correction methods which we describe in Sect. 2.

For the simulated measurement, we have sets of tracks representing the ground truth and measurement. The measurement represents a noisy subset of the ground-truth tracks, since some particle trajectories are incompletely registered by the algorithm. We shall see later in Sect. 4.1 that this can introduce a selection bias into acceleration statistics, because this process can lead to data being discarded in a manner correlated with the acceleration. In addition, the measurement contains ghost tracks, which do not closely correspond to any track in the ground-truth set. To disentangle selection bias and noise effects, we constructed a "noise-free" data set (step D). This data set was constructed by identifying all correspondences between the measurement and ground-truth tracks which were sufficiently close (within 1 px of position error over their entire lifetime). For this set of "real" tracks, we replaced the measured position with its ground-truth value. Thus, the noise-free data set represents the experimental subset of the ground-truth tracks without noise.

4 Simulation results

In this section, we apply the error correction techniques described in Sect. 2 to the numerical simulations of LPT described in Sect. 3.3. These are used to validate the use of independent measurements and spline filtering in reducing biases in acceleration moments and probability density functions due to filtering and noise.

4.1 Acceleration moments

The statistics of Lagrangian acceleration are remarkably sensitive to measurement noise, filter scale, and selection biases. Figure 4a, b shows the dependence of the measured acceleration variance upon the filter time-scale $\tau_f = 1/f_{\mathrm{ENBW}}$ for the Gaussian and spline filters, respectively. The variance is calculated for the ground truth, noise-free, and simulated measurement data sets. In addition, we plot the acceleration variance for the measurement as corrected by (3). The reference value for the acceleration variance, as obtained from the ground-truth acceleration sampled on the particles, is

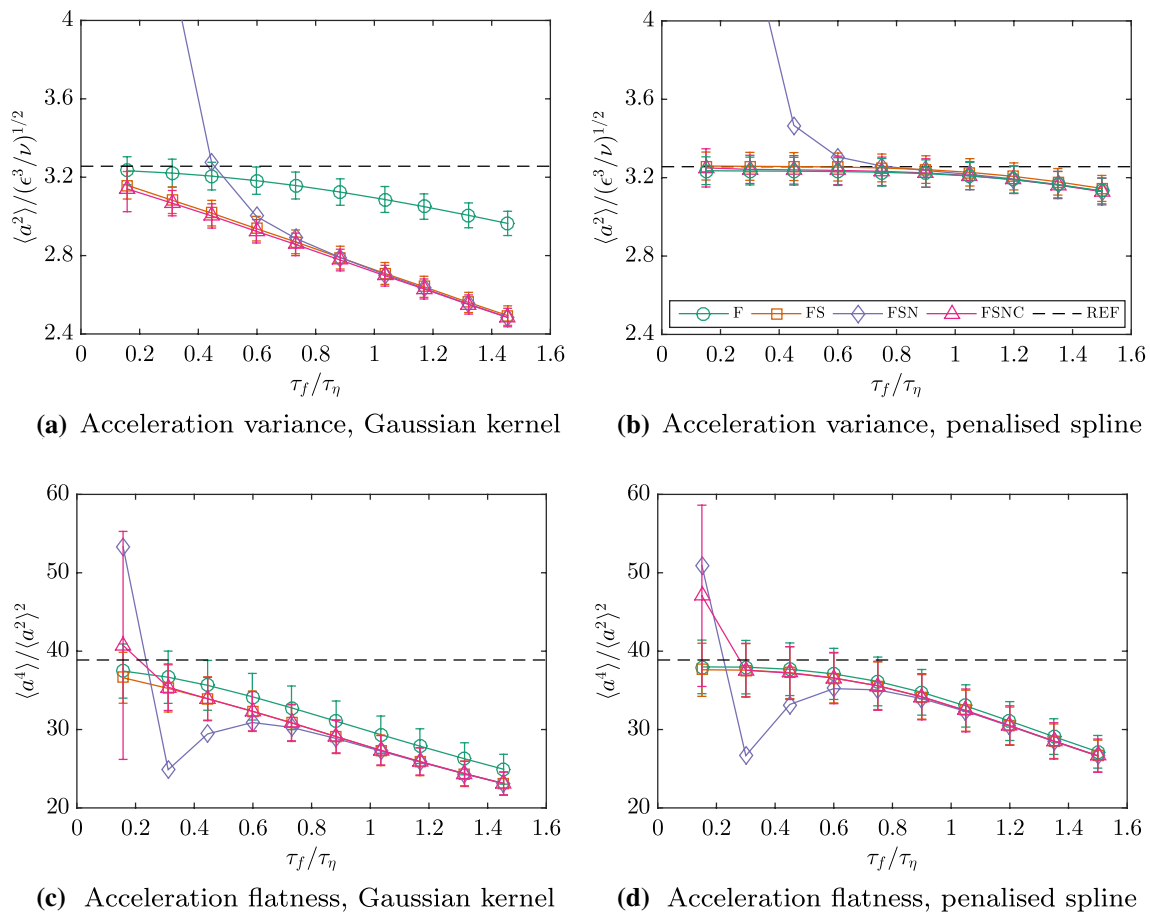indicated with the dashed line. We note the mean acceleration is negligible in all the cases.

By comparing different curves, we can isolate the effect of different error sources (filtering, selection bias, and noise). The effect of filtering is isolated by comparing the two filtered, ground-truth data sets to the reference value. The ensemble in these three cases is identical: only the manner of filtering differs. Both spline and Gaussian filters have a modest effect ($\leq 5\%$) upon acceleration variance for a filter scale below $1\tau_\eta$. The Gaussian filter attenuates the signal more than the spline: this may be attributed to the sharper spectral cutoff of the spline filter.

A more pronounced difference is seen between the two filters as applied to the noise-free data sets, shown in orange, where selection bias effects are introduced. Here, the Gaussian filter exhibits a strong dependence on scale, whilst the spline-filtered data does not. Since the Gaussian filter implementation rejects data near track ends and interpolated locations, the sampled ensemble is reduced as the filter support is increased. This introduces a selection bias by under-representing faster particles, which tend to have shorter tracks and are correlated with larger accelerations (Sawford et al. 2003; Voth et al. 1998). Omission of these data results in the underestimation of the acceleration variance.

Unsurprisingly, measurement noise significantly increases the measured acceleration variance when small filter scales are used, as evidenced by the purple curves for the synthetic measurement data set. At the smallest filter scales, the noise power exceeds the signal by a factor of ten. Due to the noise, the uncorrected measurement shows no range where the result is insensitive to the choice of scale. When noise is accounted for by applying the correction in (3), there is remarkable agreement (within 0.6%) between the corrected data (magenta) and noise-free data (green). The price paid for noise correction is statistical confidence: as the signal-to-noise ratio is reduced by decreasing the filter scale, the statistical sampling error increases.

The combined effects of filtering, selection bias, and noise are more significant when considering the higher-order moments of acceleration. Figure 4c, d shows the filter and scale dependence of the acceleration flatness for the Gaussian and spline filters, respectively. Examination of the filtered ground-truth data set shows that the acceleration flatness has a stronger dependence on filter scale when a Gaussian filter is used. Similarly, the noise-free case shows a much stronger dependence upon filter scale when a Gaussian filter is used in comparison to the smoothing spline.

The introduction of measurement noise significantly affects the measured flatness. At small filter scales, noise dominates the signal and we effectively measure the flatness factor of the noise. At larger filter scales, the filtered acceleration signal dominates, but the filtering has a strong effect on the signal's flatness factor. Thus, with uncorrected data, no range of

(a) Acceleration variance, Gaussian kernel

(b) Acceleration variance, penalised spline

(c) Acceleration flatness, Gaussian kernel

(d) Acceleration flatness, penalised spline

Fig. 4 Acceleration variance (**a**, **b**) and flatness (**c**, **d**) as a function of filter scale and type: Gaussian (**a**, **c**) and penalised spline (**b**, **d**). Black dashed line shows the reference value evaluated from DNS. Markers indicate different combinations of bias effects: ● filtering (F) ■ filtering and selection bias (FS) ◆ filtering, selection bias, and noise (FSN) ▲ filtering, selection bias, noise, and correction (FSNC). Comparisons between curves allow different bias effects to be isolated. Error bars show bootstrapped 95% statistical confidence intervals

filter scales is observed where the result is invariant to scale. Remarkably, when the noise correction (3) is applied, excellent agreement (well within statistical confidence) with the noise-free case is observed. Importantly, we observe a range of filter scales where the flatness has only a weak dependence on scale. This means that it is possible to conclude, on the basis of the data, whether the filter scale is sufficiently small.

At this stage, two conclusions may be drawn. First, it is most preferable to use cubic smoothing splines for filtering LPT data, as they introduce less selection bias and have an advantageously sharp spectral cutoff which filters less signal. Second, the noise-correction technique for moments described in Sect. 2.1 has been validated.

4.2 Acceleration PDF

We now consider the measurement of the acceleration PDF when data are contaminated with noise. To demonstrate this,

we consider the acceleration distribution obtained from the simulated measurement using a smoothing spline filter with $\tau_f = 0.27\tau_\eta$, which yields a signal-to-noise ratio of only 4 dB. Figure 5 shows the standardised PDF of acceleration (with and without correction) and directly from the DNS acceleration field (the reference). The presence of measurement noise overestimates the core of the PDF, but does not significantly influence the tails, which are in close agreement with the reference PDF. The deconvolution procedure provides a remarkably accurate correction to the measured PDF, as evidenced by the good agreement with the reference data over a wide range of kernel bandwidths.

The inset of Fig. 5 illustrates the trade-off which is made when selecting the kernel bandwidth h. Using a smaller bandwidth provides a more accurate estimation of the core, where the underlying data are dense, but undersmooths the tails of the PDF, whereas a larger bandwidth oversmooths the core but captures the tails more accurately. In general,
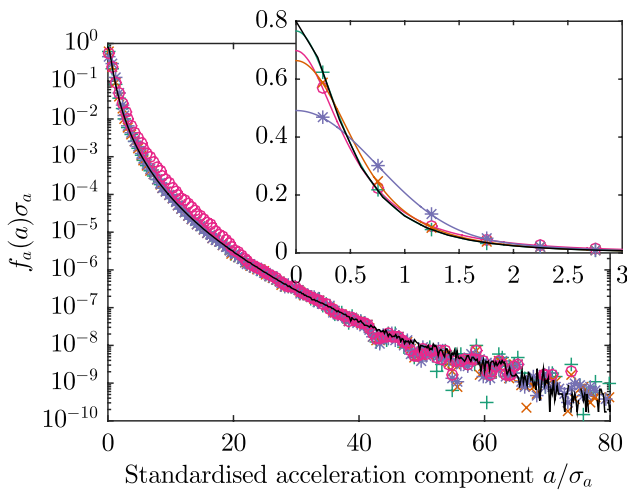
Fig. 5 Estimation and recovery of ground-truth acceleration PDF using very noisy data. The solid black line is the PDF f_a obtained from the DNS reference. Markers show: ($\circ$) the measured PDF $f_{\overline{a}}$ without correction, ($+$, $\times$, $*$) the PDF $\widetilde{f}_a$ with correction at kernel bandwidths $h/\sigma_a = 0.05, 0.10, 0.20$, respectively. Inset: the same quantities on a linear scale. Since the PDFs are symmetric, only the positive half is shown

an "optimal" bandwidth is difficult to define, since this depends upon the relative importance of statistical uncertainty and accuracy, which will depend the chosen analysis. We refer the reader to Wang and Wang (2011) for a practical review on the available methods for bandwidth selection.

One way of assessing the accuracy of the corrected PDF is to examine its moments. Based on the moment-generating properties of the characteristic function, one can show that for the kernel in (6), the second and fourth moments of the corrected PDF $\widetilde{f}_a$ are given by

$$\int \widetilde{f}_a(a)a^2 \, \mathrm{d}a = \langle \overline{a}^2 - \widehat{a}^2 \rangle + 6h^2 \tag{16}$$

and

$$\int \widetilde{f}_a(a)a^4 \, \mathrm{d}a = \langle \overline{a}^4 - \widehat{a}^4 \rangle - 6\langle \overline{a}^2 - \widehat{a}^2 \rangle(\langle \widehat{a}^2 \rangle + 6h^2) + 72h^4, \tag{17}$$

respectively. For the above choice of bandwidth and noise level, this gives a relative error of 6% in the second moment and 0.94% in the fourth moment, which are comparable to the statistical uncertainty.

5 Experimental results

In this section, we evaluate the statistics of Lagrangian acceleration for the experimental data sets described in Sect. 3.1 and make comparisons to the wider literature.

This provides a direct test of the bias correction techniques outlined in Sect. 2 over a wide range of Reynolds numbers ($R_\lambda = 109 - 504$), spatial resolution ($\eta = 1.3 - 11$ px), and effective temporal resolution ($\tau_f/\tau_\eta = 0.15 - 1.5$).

5.1 Acceleration moments

Figure 6a shows the measurements of the acceleration variance as a function of filter scale for our experimental data sets. A penalised spline filter was used. The directional dependence of this measurement is negligible ($\leq 3.1\%$ variation across cameras). Both the raw (dashed line) and noise-corrected (solid line) measurement are shown. It is readily apparent that, as the filter scale is reduced, the noise contribution dominates the measurement of acceleration variance. However, when this noise contribution is corrected for, we see that the acceleration has a weak dependence on the filter scale, which is the expected physical behaviour. When the noise level is made very large, the correction technique starts to break down. This is because, in practice, there exists a small correlation in the measurement error between cameras which becomes significant when the noise level is very large (in our data, below SNR $\simeq 0$ dB). This demonstrates that the noise-correction technique is able to accurately compensate second moments of moderately noisy experimental data.

It is worth emphasising that, when penalised spline filtering is used in conjunction with noise correction to estimate the acceleration variance, the result has a very weak dependence on filter scale. For example, at $R_\lambda = 504$ where the scale dependence is strongest, doubling the filter bandwidth from $\tau_f = 0.9\tau_\eta$ to $\tau_f = 0.45\tau_\eta$ (i.e., from $f_{\mathrm{ENBW}}\tau_\eta = 1.1$ to 2.2) increases the estimated variance by less than 3.5%, which is comparable to the statistical uncertainty. This is in stark contrast to the exponential dependence upon filter scale observed by Voth et al. (2002) and Crawford (2004). Moreover, the weak scale dependence allows a quantitative assessment of whether the filter scale is sufficiently small.

Corresponding measurements of the acceleration flatness are shown in Fig. 6b. As the filter scale is reduced, the uncorrected acceleration flatness first increases as finer scale flow features are probed, and then decreases as the noise contribution begins to dominate. By correcting for noise, we see a steady increase in the flatness factor as the filter scale is reduced. The presence of this expected physical behaviour qualitatively confirms the validity of the noise-correction approach. In contrast to the numerical simulations presented in Fig. 4d, we do not have sufficient resolution to identify clear plateau region at all but the lowest Reynolds numbers. This indicates that the experimental acceleration distribution contains contributions from very rapid motions not present in the simulations. We speculate that the difference arises due to the intermittency of the large-scale forcing in the LEM, which may occasionally generate regions of intense

(a) Variance
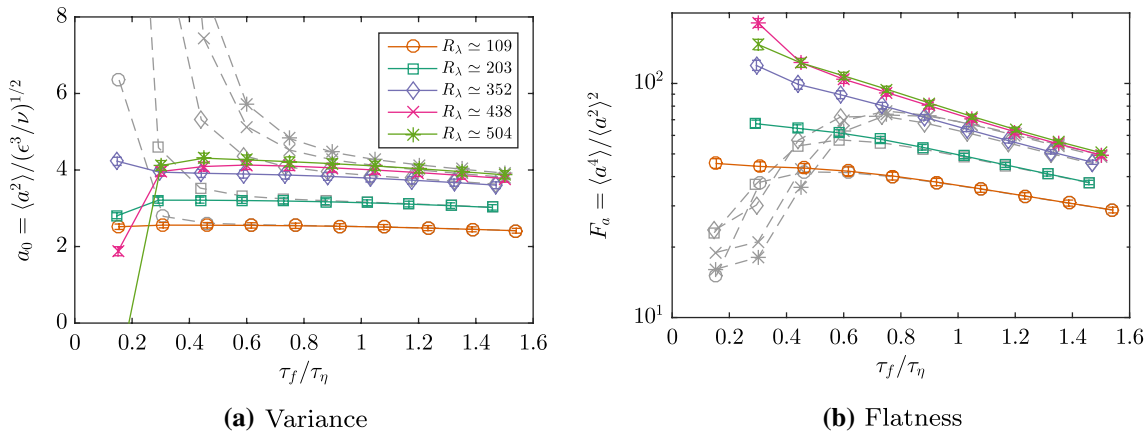
(b) Flatness

Fig. 6 Experimental measurement of acceleration **a** variance and **b** flatness factor as a function of filter scale. Solid lines show corrected moments, dashed lines show uncorrected moments. Markers

$\circ, \square, \diamond, \times, *$ show Reynolds number in ascending order. Error bars show 90% confidence intervals

turbulence where the fastest dynamics proceed on timescales below the average Kolmogorov scale.

We now compare our experimental measurements of the acceleration variance to measurements and models reported in the literature. Figure 7 shows the Heisenberg–Yaglom coefficient $a_0 = \langle a^2 \rangle / (\epsilon^3/\nu)^{1/2}$ obtained from the present experiments, simulations of homogeneous, isotropic turbulence (Gotoh and Fukayama 2001; Lalescu and Wilczek 2018b; Vedula and Yeung 1999; Yeung et al. 2006) and the multifractal model fit obtained by Sawford et al. (2003). Note that the multifractal model was obtained by (Sawford et al. 2003) by a fit to the data from Vedula and Yeung (1999) and Gotoh and Fukayama (2001). The experimental data are in reasonable agreement with the empirical fit and show a comparable degree of scatter to the available numerical data. This quantitative agreement demonstrates the ability of the experimental technique to eliminate systematic biases when measuring the acceleration variance.

5.2 Acceleration PDF

Unlike the simulation experiment, there is no ground truth available to compare the measured acceleration distribution to validate the measurement. However, we are able to test the consistency of the deconvolved PDF as a function of filter scale: when the filter scale is sufficiently small, only the tails of the distribution should change significantly. This is demonstrated in Fig. 8, which shows the corrected distribution of acceleration $\widetilde{f}_a$ at filter scales $\tau_f = 0.45 - 0.9\tau_\eta$. As the filter scale is increased, the core of the distribution remains largely unchanged, whereas the probability density in the far tails increases markedly.

It is remarkable that, even though the signal-to-noise ratio is very poor (0 dB) for $\tau_f = 0.45\tau_\eta$, the cores of the corrected PDFs are in good agreement. In this case, deconvolution
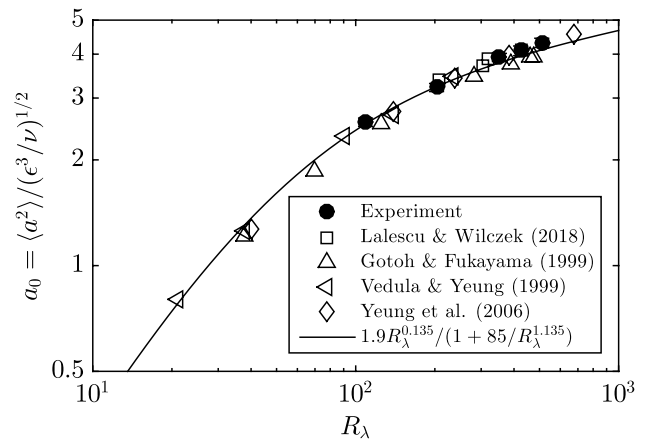

Fig. 7 Collapse of normalised acceleration variance a_0 obtained from present experimental data (for $\tau_f = 0.6\tau_\eta$) with data from simulations of homogeneous, isotropic turbulence and multifractal model fit (Sawford et al. 2003), plotted as a function of R_λ. Error bars show 90% confidence intervals

allows us to improve the temporal resolution of the measurement by a factor of approximately two. To do the same by reducing the noise level would require a fourfold improvement in position accuracy. Implementing deconvolution is, of course, substantially simpler.

6 Conclusion

We have presented methods to mitigate two key sources of bias error in LPT measurements, namely systematic errors introduced by random noise and during temporal filtering. The methods have been validated through the use of numerical simulations of LPT and demonstrated to work in practice via application to experimental LPT measurements
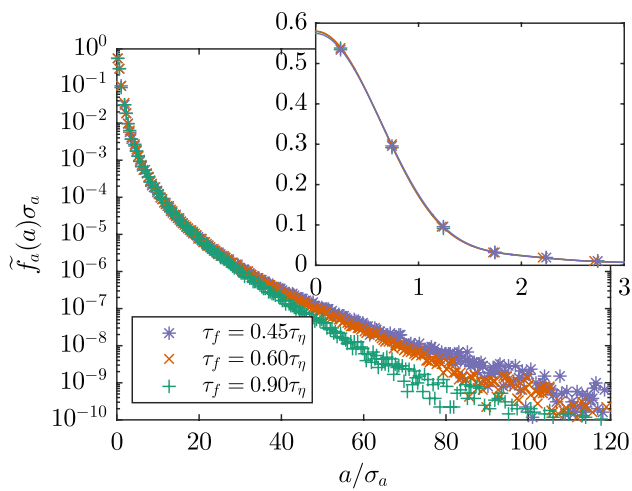
Fig. 8 Self-consistency of the corrected, experimental distribution of acceleration $\widetilde{f}_a$ at $R_\lambda = 504$, as a function of filter scale. The kernel bandwidth is $h = 0.2\sigma_a$. Inset: the same distribution on a linear scale. For clarity, only the positive half of the PDF is shown

in homogeneous, isotropic turbulence over a wide range of Reynolds numbers, and effective spatial and temporal resolutions.

We have outlined methods to correct statistical moments and probability distributions of Lagrangian quantities obtained from LPT data contaminated with random noise. The operating principle is to obtain a measurement of the noise distribution imposed on the velocity or acceleration signal using simultaneous measurements made across two or more cameras. This distribution can be used to correct the desired statistics.

The noise-correction technique is mainly limited by the increase in statistical uncertainty associated with the noise correction and the requirement that the noise remains statistically independent of the signal. The technique has potential for very general application. Examples include the measurement of velocity and Reynolds stress profiles in the near wall boundary layer, which are notoriously difficult to measure accurately (Kähler et al. 2012a).

The use of finite-impulse-response filters to measure particle acceleration can introduce a significant, systematic error in acceleration statistics. Such filters only sample the signal where there is sufficient support, which under-represents shorter, faster tracks correlated with larger accelerations. This bias is responsible for introducing a strong dependence of the measured acceleration statistics upon filter scale. Instead, we advocate the use of penalised, cubic splines to implement the numerical differentiation and filtering of tracks. When spline filtering is used, acceleration statistics are seen to display a much weaker dependence upon filter scale. This is because the acceleration may be sampled along the entire length of each track.

When spline filtering is used in combination with noise correction, a range of filter scales are observed where the acceleration statistics depend only weakly on the filter scale. This has not been achieved in the literature to date. Crucially, this allows the experimentalist to determine quantitatively whether the chosen filter bandwidth is sufficiently small. In our measurements of acceleration variance, we find that a filter bandwidth $f_{\mathrm{ENBW}}\tau_\eta \geq 1.1$ is sufficient to recover $\geq 96.5\%$ of the signal energy. The acceleration flatness requires more bandwidth, depending on the Reynolds number. The simultaneous application of these techniques allows a substantial reduction in the temporal resolution and position accuracy required to make accurate measurements of acceleration distributions in turbulent flows and Lagrangian properties of turbulent flows in general.

The generality and simplicity of the above techniques allows us to suggest the adoption of the following best practices in making Lagrangian acceleration measurements. Discrete convolution methods for numerical differentiation should be avoided due to the selection bias errors that they introduce by under-representing data near ends of tracks and in interpolated segments. Filter time-scales should be chosen in a range where they do not influence the results. Systematic errors in statistical moments and probability distributions should be quantified or corrected for via the use of simultaneous, independent measurements across independent cameras where possible.

References

Adrian RJ, Westerweel J (2011) Particle image velocimetry, 1st edn. Cambridge University Press, Cambridge

Bec J, Biferale L, Boffetta G, Celani A, Cencini M, Lanotte A, Musacchio S, Toschi F (2006) Acceleration statistics of heavy particles in turbulence. J Fluid Mech 550:349–358

Berg J, Ott S, Mann J, Lüthi B (2009) Experimental investigation of Lagrangian structure functions in turbulence. Phys Rev E 80(2):026316

Bewley GP, Saw EW, Bodenschatz E (2013) Observation of the sling effect. New J Phys 15(8):083051

Borgas MS, Sawford BL (1991) The small-scale structure of acceleration correlations and its role in the statistical theory of turbulent dispersion. J Fluid Mech 228:295–320

Crawford AM (2004) Particle tracking measurements in fully developed turbulence: water and dilute polymer solutions. Ph.D. Thesis, Cornell University

Dracos Th (1996) Particle tracking in three-dimensional space. In: Dracos Th (ed) Three-dimensional velocity and vorticity measuring and image analysis techniques. Kluwer, Dordrecht

Eilers PHC, Marx BD (1996) Flexible smoothing with B-splines and penalties. Stat Sci 11(2):89–121

Elliott DF (1987) Handbook of digital signal processing: engineering applications. Academic Press Inc, Orlando

Falkovich G, Fouxon A, Stepanov MG (2002) Acceleration of rain initiation by cloud turbulence. Nature 419(6903):151–154

Gesemann S, Huhn F, Schanz D, Schröder A (2016) From particle tracks to velocity and acceleration fields using B-splines and penalties. In: 18th International Symposium on the Appliation of Laser and Imaging Techniques to Fluid Mechanics, Lisbon, Portugal. arXiv:1510.09034

Gibert M, Xu H, Bodenschatz E (2012) Where do small, weakly inertial particles go in a turbulent flow? J Fluid Mech 698:160–167 arXiv:1002.3755

Gotoh T, Fukayama D (2001) Pressure spectrum in homogeneous turbulence. Phys Rev Lett 86(17):3775–3778

Hartley R, Zisserman A (2003) Multiple view geometry in computer vision. Cambridge University Press, Cambridge

Holzner M, Liberzon A, Nikitin N, Lüthi B, Kinzelbach W, Tsinober A (2008) A Lagrangian investigation of the small-scale features of turbulent entrainment through particle tracking and direct numerical simulation. J Fluid Mech 598:465–475

Hoyer K, Holzner M, Lüthi B, Guala M, Liberzon A, Kinzelbach W (2005) 3D scanning particle tracking velocimetry. Exp Fluids 39(5):923–934

Kähler CJ, Scharnowski S, Cierpka C (2012a) On the resolution limit of digital particle image velocimetry. Exp Fluids 52(6):1629–1639 arXiv:1405.3023

Kähler CJ, Scharnowski S, Cierpka C (2012b) On the uncertainty of digital PIV and PTV near walls. Exp Fluids 52(6):1641–1656

La Porta A, Voth GA, Crawford AM, Alexander J, Bodenschatz E (2001) Fluid particle accelerations in fully developed turbulence. Nature 409(6823):1017–1019

Lalescu CC, Wilczek M (2018a) Acceleration statistics of tracer particles in filtered turbulent fields. J Fluid Mech 847:R2

Lalescu CC, Wilczek M (2018b) How tracer particles sample the complexity of turbulence. New J Phys 20(1):013001

Maas HG, Gruen A, Papantoniou D (1993) Particle tracking velocimetry in three-dimensional flows. Exp Fluids 15(2):133–146

Malik NA, Dracos T, Papantoniou DA (1993) Particle tracking velocimetry in three-dimensional flows. Exp Fluids 15(4–5):279–294

Mann J, Ott S, Andersen JS (1999) Experimental study of relative, turbulent diffusion. Riso National Laboratory, Roskilde

Monchaux R, Bourgoin M, Cartellier A (2012) Analyzing preferential concentration and clustering of inertial particles in turbulence. Int J Multiph Flow 40:1–18

Mordant N, Crawford A, Bodenschatz E (2004) Experimental Lagrangian acceleration probability density function measurement. Phys D Nonlinear Phenom 193(1–4):245–251

Nishino K, Kasagi N, Hirata M (1989) Three-dimensional particle tracking velocimetry based on automated digital image processing. J Fluids Eng 111(4):384

Ouellette NT, Xu H, Bodenschatz E (2006) A quantitative study of three-dimensional Lagrangian particle tracking algorithms. Exp Fluids 40(2):301–313

Prasad AK, Adrian RJ, Landreth CC, Offutt PW (1992) Effect of resolution on the speed and accuracy of particle image velocimetry interrogation. Exp Fluids 13(2–3):105–116

Raffel M, Willert CE, Scarano F, Kähler CJ, Wereley ST, Kompenhans J (2018) Particle image velocimetry, 3rd edn. Springer International Publishing, Cham

Salazar JP, Collins LR (2009) Two-particle dispersion in isotropic turbulent flows. Ann Rev Fluid Mech 41(1):405–432

Sawford BL, Yeung PK, Borgas MS, Vedula P, La Porta A, Crawford AM, Bodenschatz E (2003) Conditional and unconditional acceleration statistics in turbulence. Phys Fluids 15(11):3478–3489

Schanz D, Gesemann S, Schröder A, Wieneke B, Novara M (2013) Non-uniform optical transfer functions in particle imaging: calibration and application to tomographic reconstruction. Meas Sci Technol 24(2):024009

Schanz D, Gesemann S, Schröder A (2016) Shake-The-Box: Lagrangian particle tracking at high particle image densities. Exp Fluids 57(5):70

Shaw RA (2003) Particle–turbulence interactions in atmospheric clouds. Ann Rev Fluid Mech 35(1):183–227

Stefanski LA, Carroll RJ (1990) Deconvolving kernel density estimators. Statistics 21(2):169–184

Toschi F, Bodenschatz E (2009) Lagrangian properties of particles in turbulence. Ann Rev Fluid Mech 41(1):375–404

Vedula P, Yeung PK (1999) Similarity scaling of acceleration and pressure statistics in numerical simulations of isotropic turbulence. Phys Fluids 11(5):1208–1220

Voth GA, Satyanarayan K, Bodenschatz E (1998) Lagrangian acceleration measurements at large Reynolds numbers. Phys Fluids 10(9):2268–2280

Voth GA, La Porta A, Crawford AM, Alexander J, Bodenschatz E (2002) Measurement of particle accelerations in fully developed turbulence. J Fluid Mech 469:121–160

Wang XF, Wang B (2011) Deconvolution estimation in measurement error models: the R package decon. J Stat Softw 39(10):pii:i10

Xu H (2008) Tracking Lagrangian trajectories in positionvelocity space. Meas Sci Technol 19(7):075105

Xu H, Ouellette NT, Vincenzi D, Bodenschatz E (2007) Acceleration correlations and pressure structure functions in high-reynolds number turbulence. Phys Rev Lett 99(20):204501

Yeung PK (2002) Lagrangian investigations of turbulence. Ann Rev Fluid Mech 34(1):115–142

Yeung PK, Pope SB, Lamorgese AG, Donzis DA (2006) Acceleration and dissipation statistics of numerically simulated isotropic turbulence. Phys Fluids 18(6):065103

Zimmermann R, Xu H, Gasteuil Y, Bourgoin M, Volk R, Pinton JF, Bodenschatz E (2010) The Lagrangian exploration module: an apparatus for the study of statistically homogeneous and isotropic turbulence. Rev Sci Instrum 81(5):055112