

Validation of methods for estimating velocity gradients from highly resolved particle tracking data

Daniele Certini^{1,*}, Martin Schiødt^{1,2}, Haim Abitan¹, Pratik Suchde^{3,4}, Tom Buchwald⁵, Andreas Schröder^{5,6}, Clara M. Velte¹

1: Dept. of Civil and Mechanical Engineering, Technical University of Denmark, Denmark

2: DHI, Hørsholm, Denmark

3: Department of Engineering, University of Luxembourg, Luxembourg

4: Fraunhofer Institute for Industrial Mathematics (ITWM), Germany

5: Image-Based Measurement Technology, Brandenburg University of Technology Cottbus-Senftenberg, Germany

6: Deutsches Zentrum für Luft- und Raumfahrt - DLR, Germany

*Corresponding author: dancer@dtu.dk

Keywords: LPT processing, Data assimilation, Velocity gradient estimation, Dissipation rate, Fine-scale structure of turbulence.

ABSTRACT

Velocity gradient estimation is of significant interest for turbulence theory and model development, in particular for characterising the fine scale structure of turbulence. However, correlation based methods have not generally demonstrated sufficient accuracy for measuring small scale gradients. Estimating gradients directly from particle velocity tracks can potentially provide the full velocity gradient tensor with the fidelity required for theoretical developments, since the maximum attainable information from the particle position measurements is retained. In this setting, at least for highly resolved data, measurement noise may instead become the limiting factor, as spatial differentiation amplifies noise increasingly with derivative order, even when the underlying tracks have relatively low noise levels. Herein, we compare different methods for estimating spatial velocity gradients from particle tracking data and evaluate their usability using a dataset with well defined ground truth and negligible noise levels. We compare two Eulerian reconstruction methods, Binning and VIC#-like, with two Lagrangian methods, Lagrangian Gradient Regression and the Generalized Finite Difference Method. Importantly, Binning performs local Eulerian averaging regression without enforcing continuity, whereas the VIC#-like reconstruction includes a divergence free constraint through its physics based data assimilation formulation. Therefore, the divergence residual is interpreted as an a posteriori consistency diagnostic for Binning and the Lagrangian methods, but as a constrained quantity for VIC#-like. The results show that all methods recover physically consistent gradient fields with near zero mean divergence, but differ markedly in residual scatter, sensitivity to processing parameters, and computational scalability.

1. Introduction

Spatial velocity gradients are central to turbulence theory and modeling. In particular, they are fundamental to the small-scale quantities of turbulence and not least to the small-scale universality

of turbulence assumed by Kolmogorov (1941a,b,c, 1962) and Batchelor (1953). Naturally, quantities such as the dissipation rate of turbulent kinetic energy, central to turbulence theory and modeling, are of particular interest to measure accurately. But even much more commonly measured and experimentally accessible quantities, such as the Taylor microscale, are classically connected to the small-scale behaviour or turbulence, as this scale was originally defined as the ‘dissipative scale’ of turbulence by G. I. Taylor (1935).

In experimental fluid mechanics, Lagrangian Particle Tracking (LPT) has become a prominent tool to measure particle velocity trajectories (Schröder & Schanz (2023)). By providing a Lagrangian view of velocity fields, LPT introduces both significant opportunities and notable challenges. Particle tracks, by nature, although perhaps not trivial to handle mathematically, provide exceptionally suitable information for measuring the gradient based fine-scale turbulence structure of turbulence. However, accurate spatial gradient estimation requires a sufficiently small inter-particle distance. Furthermore, although noise may be modest in a measured velocity signal, estimating its gradient signal will inevitably amplify the noise. This challenges becomes particularly prominent for higher order gradients.

Historically, data assimilation (DA) methods have been applied to analyse LPT data in an Eulerian frame of reference. Modern approaches to perform DA include VIC# (Christiansen (1973); Schneiders & Scarano (2016); Jeon et al. (2022)), FlowFit (Gesemann et al. (2016); Godbersen et al. (2024)) and neural-implicit particle advection (Zhou & Grauer (2023); Zhou et al. (2023)). Although each method offers distinct advantages depending on the application (Zhou et al. (2024)), their reliance on mapping particle data onto an Eulerian grid introduces added filtering and digitization as well as a significant computational burden. Both may be alleviated by instead adopting a Lagrangian framework for flow analysis, where quantities of interest are evaluated at particle positions.

The evaluation of spatial velocity gradients via Lagrangian flow analysis is not possible by using standard techniques such as finite differences because flow quantities are sampled at irregularly distributed particle locations.

Several alternative approaches have been proposed instead. T. Harms et al. (2023) introduced the *Lagrangian gradient regression* method, which estimates spatial velocity gradients by analyzing the local deformation of particle trajectories over time. Although this approach is appealing, because it relies solely on time-resolved particle position data, it is limited to first-order spatial gradient estimation.

Alternative approaches for estimating spatial gradients include methods based on *radial basis functions* (RBFs; Mongillo et al. (2011)), which employ RBFs as a fitting basis, and *smoothed particle hydrodynamics* (SPH; Liu & Liu (2003)), which relies on an integral representation of the target field. Although both methods are viable for gradient estimation, the need to tune multiple hyperparameters can complicate their practical use. In contrast, the *generalized finite difference method* (GFDM; Suchde (2018)) provides a more straightforward alternative and achieves comparable ac-

curacy when only low- to moderate-order gradients are required.

While Lagrangian flow analysis offers the advantage of reduced computational cost by avoiding projection onto a Eulerian grid, conservation properties are not guaranteed in all cases (Suchde et al. (2025)). The present work therefore provides a more in-depth assessment of the efficacy of Lagrangian analysis.

In order to properly test the quality of the gradient estimation, a high-quality dataset with a well defined ground truth was chosen, namely the *Second DA Challenge Case 1, Time Resolved* dataset from the ONERA flow reconstruction benchmark (*2nd LPT and DA Challenge – Case 1*, n.d.). This benchmark is based on the square duct channel flow of Pirozzoli et al. (2018). The original data consists of a Direct Numerical Simulation (DNS), with a bulk Reynolds number of 40,000. Lagrangian Particle Tracking (LPT) data is generated as a cloud of virtual particles propagated from the DNS simulation based on a four-camera arrangement. Further details can be found in Pirozzoli et al. (2018) and [here](#). At these Reynolds numbers, the square duct flow is assumed to operate in the incompressible regime, where compressibility effects are negligible in comparison to the uncertainty. Consequently, gradient accuracy can be tested and quantified immediately from continuity, i.e., the equation governing mass conservation.

The ability of the different reconstruction and differentiation methods to estimate spatial velocity gradients from high-resolution LPT data was tested on two Eulerian methods, Binning and VIC#-like, and on two Lagrangian methods, Lagrangian Gradient Regression (T. Harms et al., 2023) and GFDM (Suchde et al., 2025).

In practice, smoothed tracks are already used as input for the mentioned methods to reduce errors caused by inevitable measurement noise. Accordingly, we also employ smoothed input data. For this purpose, the particle tracks used here were preprocessed by DLR using a newly developed temporally and spatially local variant of TrackFit (Gesemann et al., 2016). The core idea is similar to the approach described in Buchwald et al. (2025), where filter parameters are adapted based on the local acceleration variance. However, the implementation used here is an advanced version which will be described in detail in a future publication.

2. Velocity gradient estimation methods

2.1. Binning

Binning (Agüí & Jiménez, 1987; Agüera et al., 2016) provides a simple and computationally inexpensive reference method for converting scattered Lagrangian particle information into an Eulerian velocity field. In contrast to data-assimilation approaches such as VIC#, binning does not explicitly impose the (incompressible) governing equations during the reconstruction. Instead, the measurement domain is divided into cubic interrogation volumes, referred to as bins, and the par-

ticles located within each bin at a given time instant are used to estimate the velocity and velocity gradients at the bin centre.

Let $\mathbf{x}_b$ denote the centre of a bin and let $\mathbf{x}_p$, $p = 1, \dots, N_b$, be the positions of the N_b particles contained within the fitting support of that bin. The local particle coordinates are defined as

$$\boldsymbol{\xi}_p = \mathbf{x}_p - \mathbf{x}_b. \quad (1)$$

The velocity component u_i is approximated locally by a second-order polynomial expansion,

$$u_i(\mathbf{x}_p) \simeq a_i + a_{ij}\xi_{p,j} + \frac{1}{2}a_{ijk}\xi_{p,j}\xi_{p,k}, \quad (2)$$

where a_i is the fitted velocity component at the bin centre, a_{ij} represents the first-order spatial velocity gradients at the bin centre, and a_{ijk} contains the second-order polynomial coefficients. The coefficients are obtained independently for each velocity component by solving a least-squares problem,

$$\mathbf{a}_i = \arg \min_{\mathbf{a}_i} \sum_{p=1}^{N_b} [u_i(\mathbf{x}_p) - \boldsymbol{\phi}(\boldsymbol{\xi}_p)^T \mathbf{a}_i]^2, \quad (3)$$

where $\boldsymbol{\phi}$ is the vector of polynomial basis functions.

The binned velocity field is then obtained from the constant coefficient of the polynomial fit,

$$u_i(\mathbf{x}_b) = a_i, \quad (4)$$

while the first-order velocity-gradient tensor follows directly from the linear coefficients,

$$\left. \frac{\partial u_i}{\partial x_j} \right|_{\mathbf{x}_b} = a_{ij}. \quad (5)$$

Although the second-order polynomial coefficients are included in the local fit, only the velocity and first-order velocity gradients at the bin centre are retained for the present analysis.

Before binning, particle velocities are obtained from the time resolved tracks using a locally adjusted TrackFit procedure.

After the local polynomial fit, the residual divergence is evaluated a posteriori. This step is particularly relevant for incompressible flows, where

$$\nabla \cdot \mathbf{u} = 0 \quad (6)$$

provides a physical consistency condition. In the present comparison, the residual divergence of the reconstructed field is therefore used as one of the primary validation metrics. The pressure field, when required, can be obtained from the reconstructed velocity field by solving a pressure

Poisson equation. However, for the present study, the main quantities of interest are the velocity gradient tensor and derivative based diagnostics.

The principal advantage of binning is its robustness and low computational cost. It provides a transparent baseline against which more sophisticated methods can be compared. Its limitations are also clear. The spatial resolution is constrained by the bin size and by the number of particles required for a stable polynomial fit. Empty or weakly populated bins can lead to missing or unreliable vectors, while excessive bin sizes smooth the velocity field and attenuate small-scale gradients. These limitations are especially important in the present work, where the goal is to evaluate velocity gradients at scales comparable to the Kolmogorov length scale. For this reason, the bin size, overlap, and minimum particle count are important processing parameters and must be considered when interpreting the resulting divergence residuals and gradient statistics.

2.2. VIC#-like

The VIC#-like method considered here is inspired by the Vortex-In-Cell family of data-assimilation approaches, including VIC+ (Schneiders & Scarano, 2016) and VIC# (Jeon et al., 2022). However, the implementation used in the present work is not the proprietary DaVis VIC# algorithm. It is a simplified MATLAB reconstruction workflow designed to provide an Eulerian, approximately divergence-constrained reference method for comparison with the Lagrangian gradient estimators.

The method first reconstructs the velocity field on a Cartesian grid from the TrackFit-preprocessed particle data. For each grid point, particles located within a prescribed spherical support radius are selected. The local particle coordinates are defined relative to the grid point, and the three velocity components are reconstructed by solving a weighted least-squares polynomial fit. In the present implementation, a first-order spatial polynomial is used. The particle weights are defined by a Gaussian function,

$$w_p = \exp \left[- \left(\frac{r_p}{\sigma} \right)^2 \right], \quad \sigma = \frac{R}{2}, \quad (7)$$

where r_p is the distance between particle p and the grid point, and R is the support radius. A small Tikhonov regularisation term is included to stabilise sparse or ill-conditioned local fits. The reconstructed velocity at the grid point is given by the constant coefficient of the local polynomial fit.

The reconstruction is performed independently at each time step. Missing grid values are filled to allow finite-difference operations on the complete Cartesian grid. The reconstructed velocity field is then projected towards a divergence-free field using a Helmholtz-like projection,

$$\mathbf{u}_{\text{proj}} = \mathbf{u} - \nabla \phi, \quad (8)$$

where ϕ is obtained from

$$\nabla^2 \phi = \nabla \cdot \mathbf{u}. \quad (9)$$

The Poisson problem is solved with approximate homogeneous Neumann boundary conditions. This projection represents the main incompressibility-inspired constraint in the present MATLAB implementation. In contrast to full VIC# or FlowFit-type data-assimilation methods, the present implementation does not assimilate material acceleration, does not reconstruct pressure, and does not solve a global vorticity-based optimisation problem.

After projection, the spatial velocity-gradient tensor is computed on the Cartesian grid using finite differences,

$$\nabla \mathbf{u} = \begin{bmatrix} \partial u / \partial x & \partial u / \partial y & \partial u / \partial z \\ \partial v / \partial x & \partial v / \partial y & \partial v / \partial z \\ \partial w / \partial x & \partial w / \partial y & \partial w / \partial z \end{bmatrix}. \quad (10)$$

Central differences are used at interior grid points, while one-sided differences are used at the domain boundaries. The divergence, vorticity magnitude, and Frobenius norm of the velocity-gradient tensor are then evaluated from these gradients. As for the other methods, the normalised divergence residual,

$$\frac{\nabla \cdot \mathbf{u}}{U_\infty / D}, \quad (11)$$

is used as a physical consistency diagnostic for the reconstructed gradient field.

The VIC#-like method therefore differs from Binning in two main respects. First, it uses a spherical support and Gaussian-weighted local reconstruction around Cartesian grid points. Second, the reconstructed gridded field is projected towards incompressibility before the velocity gradients are evaluated. The method should therefore be interpreted as a simplified, VIC#-inspired Eulerian reconstruction and projection workflow, rather than as a full implementation of the VIC# algorithm.

2.3. Lagrangian Gradient Regression

Lagrangian Gradient Regression (LGR) estimates the local velocity-gradient tensor directly from time-resolved particle trajectories, without first reconstructing the velocity field on an Eulerian grid (T. Harms et al., 2023). In the present implementation, LGR is used as a short-time, instantaneous gradient estimator between consecutive particle frames, rather than for finite-time coherent-structure diagnostics such as Finite-time Lyapunov exponents (FTLE) or Lagrangian-averaged vorticity deviation (LAVD) (T. D. Harms et al., 2024).

Following the mathematical formulation of T. D. Harms et al. (2024), LGR begins by considering the trajectory of a passive tracer governed by

$$\frac{d\mathbf{x}}{dt} = \mathbf{u}(\mathbf{x}, t), \quad (12)$$

and the flow map

$$\mathbf{F}_{t_0}^t(\mathbf{x}_0) = \mathbf{x}(t; t_0, \mathbf{x}_0) \quad (13)$$

maps the initial tracer position $\mathbf{x}_0$ to its position at time t . The local deformation of an infinitesimal material neighbourhood is described by the flow-map Jacobian,

$$\mathbf{G} = \mathbf{D}\mathbf{F}_{t_0}^t = \frac{\partial \mathbf{F}_{t_0}^t}{\partial \mathbf{x}_0}. \quad (14)$$

For a sufficiently small time increment $\Delta t = t - t_0$, the flow-map Jacobian is related to the instantaneous velocity-gradient tensor by

$$\mathbf{G} \simeq \mathbf{I} + \nabla \mathbf{u} \Delta t, \quad \nabla \mathbf{u} \simeq \frac{\mathbf{G} - \mathbf{I}}{\Delta t}. \quad (15)$$

In practice, the deformation tensor is estimated from the relative motion of neighbouring particles that are present in both consecutive frames. For a central particle $\mathbf{x}_p$ and n neighbouring particles $\mathbf{x}_i$, the relative positions at the two times are assembled as

$$\mathbf{X}_0 = [\mathbf{x}_1(t_0) - \mathbf{x}_p(t_0), \dots, \mathbf{x}_n(t_0) - \mathbf{x}_p(t_0)], \quad (16)$$

$$\mathbf{X}_{\Delta t} = [\mathbf{x}_1(t_0 + \Delta t) - \mathbf{x}_p(t_0 + \Delta t), \dots, \mathbf{x}_n(t_0 + \Delta t) - \mathbf{x}_p(t_0 + \Delta t)]. \quad (17)$$

The best linear map between the initial and deformed neighbourhoods is then obtained through a weighted, regularised least-squares regression,

$$\mathbf{G} = \mathbf{X}_{\Delta t} \mathbf{K} \mathbf{X}_0^T (\mathbf{X}_0 \mathbf{K} \mathbf{X}_0^T + \gamma n \mathbf{I})^{-1}, \quad (18)$$

where $\mathbf{K}$ contains the neighbour weights and γ is a regularisation parameter. In the present implementation, neighbours are selected within a prescribed support radius and Gaussian distance-based weights are used. A full-rank neighbourhood is required, and ill-conditioned local regressions are rejected.

Finally, substituting Eq. (18) into Eq. (15) gives the local estimate of $\partial u_i / \partial x_j$ at the central particle position. The divergence, vorticity magnitude, and Frobenius norm of the velocity-gradient tensor are then evaluated a posteriori. LGR is therefore attractive for Lagrangian particle tracking data because it avoids interpolation to a grid and numerical differentiation of noisy velocity fields (T. Harms et al., 2023; T. D. Harms et al., 2024). Its resolution and robustness are controlled by the size and geometry of the particle neighbourhood: larger supports improve conditioning but act as a spatial low-pass filter on the estimated gradients.

2.4. Generalized Finite Difference Method

The Generalized Finite Difference Method (GFDM) provides a way to evaluate spatial derivatives directly on irregular point sets, making it well suited for LPT data where samples are available only at scattered particle locations. In the classical irregular-grid finite-difference literature, derivative stencils are constructed from local neighbourhoods and derived from Taylor-series expansions, enabling finite-difference-like operators without a structured mesh (Liszka & Orkisz, 1980). In the present work, we use a particle-centred GFDM implementation based on the direct differential-operator formulation discussed by Suchde (2018).

In the present implementation, each particle is treated in turn as the central evaluation point. A local stencil of neighbouring particles is selected around this central particle, and the central particle is included as the first stencil point. For a central particle at position $\mathbf{x}_p$, the local coordinate offsets of the stencil particles are defined as

$$\delta\mathbf{x}_{pq} = \mathbf{x}_q - \mathbf{x}_p, \quad q = 1, \dots, N_s, \quad (19)$$

where N_s is the number of particles in the local stencil. In the active implementation, neighbours are selected using a radius-based search, although a k -nearest-neighbour option is also available. The local operator construction uses quadratic polynomial consistency in three dimensions and therefore requires a sufficiently populated and well-conditioned stencil.

The distance-dependent stencil weights are defined as

$$W_q = \exp \left[-\beta \left(\frac{r_q}{h} \right)^2 \right], \quad (20)$$

where $r_q = \|\delta\mathbf{x}_{pq}\|$, h is the local smoothing length, and β is the Gaussian shape parameter. In the radius-based mode used here, h is taken as the support radius. Stencils that contain too few particles, do not satisfy the required rank condition, or exceed the prescribed condition-number threshold are rejected.

For each valid stencil, Suchde-style first-derivative operators are computed for the three coordinate directions. Denoting these operators by $\mathbf{D}_x$, $\mathbf{D}_y$, and $\mathbf{D}_z$, the local velocity-gradient components are obtained by applying the operators directly to the velocity values in the stencil,

$$\left. \frac{\partial u_i}{\partial x_j} \right|_{\mathbf{x}_p} \simeq \mathbf{D}_j^T \mathbf{u}_{i,s}, \quad (21)$$

where $\mathbf{u}_{i,s}$ contains the i -th velocity component of the stencil particles and $j = x, y, z$. The resulting estimate is therefore particle-centred and does not require interpolation to an Eulerian grid.

After the velocity-gradient tensor has been estimated, the divergence,

$$\nabla \cdot \mathbf{u} = \frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} + \frac{\partial w}{\partial z}, \quad (22)$$

the vorticity magnitude, and the Frobenius norm of the velocity-gradient tensor are evaluated a posteriori. As for LGR, GFDM does not impose incompressibility during the gradient estimation. Consequently, the normalised divergence residual is used as a diagnostic of the local consistency of the estimated gradient field.

Near boundaries, local stencils may become geometrically unbalanced because neighbouring particles are available predominantly on one side of the central particle. This can introduce systematic derivative bias and makes the accuracy of GFDM sensitive to the support size, particle density, and boundary treatment. These sensitivities are particularly important in the present comparison, where GFDM is used as a direct Lagrangian gradient estimator on irregular particle data.

2.5. FlowFit

FlowFit reconstructs instantaneous three-dimensional flow fields by representing the velocity field using a volumetric B-spline basis. The spline coefficients are fitted to particle trajectory data obtained from the TrackFit algorithm of Gesemann et al. (2016). Owing to the analytical properties of B-splines, the reconstructed velocity field and its spatial derivatives can be evaluated efficiently at arbitrary locations within the domain, facilitating direct computation of velocity gradients at selected points.

The spline coefficients are determined through the minimization of a global objective function. This function includes residuals of the particle tracks at their measurement locations as well as residuals associated with the incompressible Navier–Stokes equations. The optimisation is performed using a gradient-based method, with the coefficients initialized from a coarse first-order approximation.

In its current implementation, FlowFit operates as a purely three-dimensional method applied independently at each time step. It uses particle velocities and material derivatives provided by TrackFit as input data. The latest version, FlowFit3 (Godbersen et al., 2024), introduces an improved treatment of mass conservation. Rather than enforcing continuity through a penalty term in the objective function, divergence-free velocity fields are imposed as a hard constraint. This is achieved by restricting the optimizer to a divergence-free subspace, thereby preventing the exploration of non-physical states and accelerating convergence.

Two operational modes are available: a linear variant (FF-linear) and a nonlinear variant (FF-nonlinear).

In FF-linear, mass continuity is enforced and the optimisation is driven solely by particle velocity data. The objective function combines the LPT data residual with a regularisation term that penalises excessive curvature of the velocity field. This formulation is computationally efficient but does not reconstruct the pressure field.

In contrast, FF-nonlinear incorporates the momentum equation and particle acceleration data from TrackFit. The flow is parameterised exclusively through velocity coefficients, while the pressure field is obtained via a fast Fourier transform–based Poisson solver. By incorporating additional physical constraints and acceleration information, FF-nonlinear improves the accuracy of the reconstructed velocity and gradient fields and enables pressure recovery. This enhanced fidelity, however, comes at increased computational cost relative to FF-linear.

Both variants employ a limited-memory Broyden–Fletcher–Goldfarb–Shanno (L-BFGS) algorithm for optimization, with termination based on stagnation of the objective function. In the present study, parameters are kept largely consistent across cases. A B-spline grid resolution corresponding to 0.2 particles per cell (ppc) is used. For cases with lower seeding density, the grid is coarsened accordingly to reduce computational effort. This adaptive coarsening is disabled for the homogeneous isotropic turbulence case, where a fixed grid corresponding to 0.2 ppc in the densest configuration is retained for all realizations. A curvature penalization factor of 0.03 is applied in all cases, with minor adjustments introduced in a few extreme conditions following qualitative inspection. FlowFit is retained in the present manuscript to provide methodological context within the broader class of physics-informed Eulerian data-assimilation approaches; however, FlowFit results are not included in the quantitative comparison presented below, which is restricted to Binning, VIC#-like, GFDM, and LGR.

3. Results: Comparison of velocity gradient diagnostics

As an independent benchmark of the GFDM implementation, the method was evaluated using the ONERA Second Data Assimilation Challenge Case 1 Time Resolved dataset. The submitted algorithm, GFDM was assessed by the ONERA online evaluation platform over four particle image densities, $ppp = 0.001, 0.010, 0.050$, and 0.200 . The benchmark provides error statistics for the reconstructed velocity and vorticity fields as a function of seeding density, together with spatial error slices for a selected case.

Figure 1 shows that the velocity error decreases systematically with increasing particle density. The velocity bias is largest at the lowest seeding density and becomes small for $ppp \geq 0.05$, while the random velocity error decreases monotonically. The vorticity error is more demanding. Although the random vorticity error also decreases with increasing ppp , the GFDM results remain above the selected reference methods across the tested densities. This behaviour is consistent with the expected sensitivity of derivative based quantities to particle spacing, local stencil conditioning, and boundary effects.

The spatial error slices at $ppp = 0.05$, presented in Figure 2, further illustrate this behaviour. The streamwise velocity error is generally low over most of the plane, with larger values concentrated close to the lower and lateral boundaries. The vorticity magnitude error is more spatially struc-




Figure 1. ONERA benchmark results for the Second Data Assimilation Challenge Case 1 time-resolved dataset. The panels show velocity bias error, velocity random error, vorticity bias error, and vorticity random error as functions of ppp . Line legend GFDM, — ■; NIPA, - ■; DLR-PINN, - ■; nudging, - ■.




Figure 2. Spatial error slices from the ONERA Second Data Assimilation Challenge Case 1 time-resolved benchmark for the GFDM submission at $ppp = 0.05$. Left: streamwise velocity error magnitude, $|\epsilon_{v_x}|$, normalised by U_∞ . Right: vorticity-magnitude error, $|\epsilon_\omega|$, normalised by U_∞/D . The slices are shown in the Y - Z plane and highlight that the velocity error remains comparatively small over most of the domain, whereas the vorticity error exhibits stronger spatial structure and local amplification, particularly near the lower boundary.

tured and exhibits stronger local amplification near the same regions. This confirms that the main limitation of the present implementation is not the reconstruction of velocity itself, but the robustness of the local derivative estimates, especially where the particle support becomes asymmetric or poorly conditioned.

Overall, the ONERA benchmark supports the use of GFDM as a promising direct Lagrangian gradient estimator, but also shows that the method requires careful control of support size, particle density, stencil conditioning, and boundary treatment.

The normalised divergence residual, $(\nabla \cdot \mathbf{u})/(U_\infty/D)$, was used as a diagnostic of the internal consistency of the estimated velocity-gradient fields. For an incompressible flow, the expected value of this quantity is zero; therefore, the centring of the probability density function indicates possible systematic bias, while its width quantifies the scatter of the gradient estimates. The corresponding one-dimensional PDFs are shown separately for the Eulerian methods in Figure 3 and for the Lagrangian methods in Figure 4. The mean values and standard deviations are reported in Table 1.

For the Eulerian methods, both Binning and VIC#-like produce divergence-residual distributions centred close to zero over the range of particle image densities considered. This indicates that neither method introduces a strong systematic divergence bias. The main difference lies in the width of the residual distributions. For the processing parameters used here, Binning gives smaller standard deviations than VIC#-like at all ppp values. The corresponding values are 6.919×10^{-2} and 2.817×10^{-1} at $ppp = 0.001$, 1.468×10^{-1} and 2.875×10^{-1} at $ppp = 0.010$, 9.418×10^{-2} and 1.622×10^{-1} at $ppp = 0.050$, and 7.383×10^{-2} and 1.385×10^{-1} at $ppp = 0.200$, for Binning and

ppp	Method	Mean	Standard deviation
0.001	Binning*	-3.099×10^{-3}	6.919×10^{-2}
0.001	VIC#-like	1.014×10^{-2}	2.817×10^{-1}
0.001	GFDM	2.650×10^{-2}	3.922×10^{-1}
0.001	LGR	2.054×10^{-2}	3.522×10^{-1}
0.010	Binning	2.087×10^{-3}	1.468×10^{-1}
0.010	VIC#-like	3.932×10^{-3}	2.875×10^{-1}
0.010	GFDM	7.295×10^{-3}	4.287×10^{-1}
0.010	LGR	1.063×10^{-3}	4.485×10^{-1}
0.050	Binning	-4.110×10^{-4}	9.418×10^{-2}
0.050	VIC#-like	2.691×10^{-3}	1.622×10^{-1}
0.050	GFDM	1.154×10^{-3}	4.707×10^{-1}
0.050	LGR	-1.237×10^{-3}	5.362×10^{-1}
0.200	Binning	9.757×10^{-4}	7.383×10^{-2}
0.200	VIC#-like	-5.670×10^{-4}	1.385×10^{-1}
0.200	GFDM	3.891×10^{-3}	5.672×10^{-2}
0.200	LGR [†]	–	–

Table 1. Mean and standard deviation of the normalised divergence residual, $(\nabla \cdot \mathbf{u})/(U_\infty/D)$, for the four gradient-estimation methods at the four ppp values studied.

*The Binning result at ppp = 0.001 was computed using a larger support volume, 50 mm instead of the nominal 20 mm, to obtain a stable estimate at the lowest particle density.

[†]The LGR result at ppp = 0.200 was not computed because of computational limitations.




Figure 3. Probability density functions of the normalised divergence residual, $(\nabla \cdot \mathbf{u})/(U_\infty/D)$, for the Eulerian methods Binning and VIC#-like at four particle image densities. Rows correspond, from top to bottom, to particle image densities $\text{ppp} = 0.001, 0.010, 0.050$, and 0.200 , respectively. The dashed vertical line indicates the incompressible limit, $\nabla \cdot \mathbf{u} = 0$. The Binning result at $\text{ppp} = 0.001$ was computed using a larger support volume, 50 mm instead of the nominal 20 mm, to obtain a stable estimate at the lowest particle density.




Figure 4. Probability density functions of the normalised divergence residual, $(\nabla \cdot \mathbf{u}) / (U_\infty / D)$, for the Lagrangian gradient estimation methods GFDM and LGR. Rows correspond, from top to bottom, to particle image densities $\text{ppp} = 0.001, 0.010, 0.050$, and 0.200 , respectively. The dashed vertical line indicates the incompressible limit, $\nabla \cdot \mathbf{u} = 0$. The LGR result at $\text{ppp} = 0.200$ was not computed, because of computational limitations.

VIC#-like respectively. The Binning case at $\text{ppp} = 0.001$ should, however, be interpreted separately, since it was computed using a larger support volume to stabilise the estimate at the lowest particle density and is therefore not directly equivalent in spatial resolution to the other Binning cases.

The joint PDFs of the continuity-balance terms provide a complementary diagnostic of the same behaviour. Figure 5 compares $-(\partial u / \partial x) / (U_\infty / D)$ with $[(\partial v / \partial y) + (\partial w / \partial z)] / (U_\infty / D)$ for the two Eulerian methods. In both cases, the probability density is predominantly aligned with the diagonal line corresponding to the incompressible balance. This confirms that the two sides of the continuity relation are correlated and that the residual divergence mainly reflects scatter about this balance rather than a large systematic offset. The Binning distributions are comparatively compact, whereas the VIC#-like distributions show a broader spread around the diagonal, consistent with the standard deviations reported in Table 1.

For the Lagrangian methods, both GFDM and LGR also yield divergence-residual PDFs approximately centred around zero, as shown in Figure 4. The mean values remain small compared with the distribution widths, indicating that the dominant error is local scatter in the estimated gradients rather than a systematic divergence offset. At $\text{ppp} = 0.001$, LGR gives a slightly smaller standard deviation than GFDM, while GFDM gives slightly narrower distributions at $\text{ppp} = 0.010$ and 0.050 . The LGR result at $\text{ppp} = 0.200$ was not computed because of computational limitations, whereas GFDM shows a marked reduction in spread at this highest particle image density.

The Lagrangian continuity-balance joint PDFs are shown in Figure 6. Both GFDM and LGR produce distributions that remain organised around the incompressible diagonal, demonstrating that the estimated gradient components retain the expected physical correlation. However, compared with the Eulerian cases, the Lagrangian distributions are broader for most particle image densities, reflecting the challenge of estimating gradients directly from irregular particle data without a global incompressibility constraint. Overall, all four methods produce approximately zero-centred divergence residuals, while the principal differences are expressed in the spread of the residual and in the compactness of the joint PDFs around the incompressible continuity-balance line.




Figure 5. Joint probability density functions of the continuity-balance terms, $-(\partial u / \partial x) / (U_\infty / D)$ and $[(\partial v / \partial y) + (\partial w / \partial z)] / (U_\infty / D)$, for the Eulerian methods Binning and VIC#-like. Rows correspond, from top to bottom, to particle image densities $\text{ppp} = 0.001, 0.010, 0.050$, and 0.200 , respectively. The black diagonal line indicates the incompressible continuity balance, $-(\partial u / \partial x) = (\partial v / \partial y) + (\partial w / \partial z)$. The colour scale represents $\log_{10}(\text{PDF})$ and is common to all panels.




Figure 6. Joint probability density functions of the continuity-balance terms, $-(\partial u / \partial x) / (U_\infty / D)$ and $[(\partial v / \partial y) + (\partial w / \partial z)] / (U_\infty / D)$, for the Lagrangian methods GFD and LGR. Rows correspond, from top to bottom, to particle image densities $ppp = 0.001, 0.010, 0.050$, and 0.200 , respectively. The black diagonal line indicates the incompressible continuity balance, $-(\partial u / \partial x) = (\partial v / \partial y) + (\partial w / \partial z)$. The colour scale represents $\log_{10}(\text{PDF})$ and is common to all computed panels. The LGR case at $ppp = 0.200$ was not computed, because of computational limitations.

4. Conclusions

The four methods, two Eulerian, Binning and VIC#-like, and two Lagrangian, GFDM and LGR, were implemented and evaluated using the normalised divergence residual and the continuity-balance joint PDF. However, the processing parameters were not varied systematically. Therefore, the results should not be presented as a conclusive ranking of Binning, VIC#-like, GFDM and LGR, but as an initial assessment of their behaviour under the selected processing conditions.

The divergence residual PDFs are generally centred close to zero, as quantified in Table 1, and the continuity-balance joint PDFs remain organised around the incompressible diagonal. This indicates that all methods preserve the expected physical relation between the velocity-gradient components at least in a statistical sense. This is an important positive result, because it shows that the implemented pipelines produce physically interpretable gradient fields.

The mean divergence residuals reported in Table 1 are small compared with their standard deviations. Therefore, the dominant distinction between methods lies in the width of the residual distributions. This width reflects the local robustness of the gradient estimates under the selected parameters, rather than a systematic violation of incompressibility.

For the tested settings, Binning gives smaller standard deviations than VIC#-like and the Lagrangian methods over most of the tested ppp range. This outcome is strongly linked to the selected support volume and to the smoothing implicit in the binning operation. In particular, the $\text{ppp} = 0.001$ case required a larger support volume to include enough particles, so the corresponding residual distribution partly reflects reduced spatial resolution. Binning should therefore be described as the most stable method for this diagnostic under the present settings, not necessarily as the most accurate method for gradient recovery.

The VIC#-like fields produce divergence residual PDFs centred close to zero and joint PDFs aligned with the incompressible continuity balance. The divergence free constraint is embedded in the physics-based data assimilation formulation used to reconstruct the Eulerian velocity field from the particle data. The broader residual distributions observed may depend on reconstruction parameters, grid resolution, regularisation, and convergence settings. Without a dedicated sensitivity study, the present comparison should be interpreted as representative of the implemented configuration rather than of the full potential of VIC#.

The Lagrangian methods are particularly relevant because they estimate gradients directly from particle data, avoiding the intermediate projection onto an Eulerian grid. This is the central methodological direction of the study. Nevertheless, both GFDM and LGR depend on parameters that were not systematically optimised for this study. For GFDM, these include the support radius, the minimum number of neighbours, the weighting function, and conditioning criteria. For LGR, they include the temporal regression interval, neighbour selection, and regression settings. Therefore,

the present results should be used to motivate a more rigorous Lagrangian method comparison rather than to close it. The absence of the LGR result at $\text{ppp} = 0.200$ also shows that computational feasibility becomes a practical limitation at high particle density.

Overall, the present comparison should be regarded as a first implementation-level assessment of four velocity-gradient estimation strategies rather than as a definitive optimisation study. The normalised divergence residuals and continuity balance joint PDFs show that all methods produce gradient fields that are broadly consistent with incompressibility, with mean residuals close to zero and probability densities organised around the expected continuity balance; however, for VIC#-like this consistency is partly expected because incompressibility enters the reconstruction itself. The principal differences are therefore associated with the scatter of the residuals and with the practical robustness of each method under the selected processing parameters. Under these conditions, Binning yields the most compact divergence residual distributions, although this result must be interpreted in light of the spatial averaging associated with the support volume. VIC#-like, GFDM and LGR also retain the expected physical correlation between gradient components, but display broader residual distributions for the parameters used here.

Among the Lagrangian approaches, GFDM is a particularly promising route for direct gradient estimation because it evaluates velocity gradients directly at particle locations, without requiring projection onto an Eulerian grid. This makes it attractive for highly resolved particle-tracking measurements, where preserving the native irregular sampling may reduce interpolation and filtering effects. However, the present results also show that GFDM should not be regarded as parameter-free. Its accuracy and robustness depend strongly on support-size selection, local particle density, stencil conditioning, and boundary treatment. Excessively small supports may lead to poorly conditioned stencils, whereas excessively large supports improve robustness at the cost of spatial resolution.

A future study should systematically vary the main GFDM and LGR input parameters and quantify their influence on divergence residuals, continuity-balance joint PDFs, vorticity, strain rate, dissipation-related quantities, and computational cost. This would allow the comparison to move from an implementation-level diagnostic assessment to a rigorous methodological validation before drawing general conclusions on the relative accuracy and applicability of each method to highly resolved PTV measurements.

A selection of the evaluated methods will in coming work be used towards high-resolution PTV measurements in a round turbulent jet, initiated in Certini et al. (2026). These measurements are unique in many respects; the Reynolds number is high ($\sim 10^5$) for laboratory scale tests, while still able to resolve both the dissipation rate of the turbulent kinetic energy and even the full dissipation transport balance, including dissipation of the dissipation. Hence, filtering effects are expected to be negligible while noise on the particle tracks, amplifying with higher order derivatives, is expected to constitute the main bottleneck. The GFDM method inherently smooths the noise out

in the Taylor expansion where the different order spatial velocity gradients are estimated, based on a neighbourhood particle cluster motion. A larger cluster increasingly smooths the gradient estimates at the cost of spatial resolution, so a trade-off at a measured high particle track density is used. The implementation of the GFDM to these measurement data is reported in Certini et al. (2026).

5. Acknowledgements

CMV and MS acknowledge the support from the European Research Council: This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 803419). CMV, HA and DC acknowledge the support from the Poul Due Jensen Foundation grant No 2018-039: Financial support from the Poul Due Jensen Foundation (Grundfos Foundation) for this research is gratefully acknowledged.

References

- 2nd *lpt and da challenge – case 1*. (n.d.). (https://onera-flow-benchmarks.fr/en/Second_DA_1_Datasets)
- Agüera, N., Cafiero, G., Astarita, T., & Discetti, S. (2016). Ensemble 3d ptv for high resolution turbulent statistics. *Measurement Science and Technology*, 27(12), 124011. doi: 10.1088/0957-0233/27/12/124011
- Agüí, J. C., & Jiménez, J. (1987). On the performance of particle tracking. *Journal of Fluid Mechanics*, 185, 447–468. doi: 10.1017/S0022112087003252
- Batchelor, G. K. (1953). *The Theory of Homogeneous Turbulence*. Cambridge U. Press.
- Buchwald, T., Schanz, D., & Schröder, A. (2025). Accurate determination of acceleration statistics in a turbulent jet using locally adapted TrackFit. In *Proceedings of the 16th International Symposium on Particle Image Velocimetry, Tokyo, Japan*.
- Certini, D., Abitan, H., Ribergaard, S. L., Schiødt, M., & Velte, C. M. (2026). Measurements of the full instantaneous dissipation rate & dissipation transport in a high-Reynolds number turbulent round jet. In *Proceedings of the 22nd International Symposium on Applications of Laser and Imaging Techniques to Fluid Mechanics, Lisbon, Portugal, 29 June - 2 July 2026*.
- Christiansen, I. (1973). Numerical simulation of hydrodynamics by the method of point vortices. *Journal of Computational Physics*, 13(3), 363–379.

- Gesemann, S., Huhn, F., Schanz, D., & Schröder, A. (2016). From noisy particle tracks to velocity, acceleration and pressure fields using B-splines and penalties. In *Proceedings of the 18th International Symposium on Applications of Laser and Imaging Techniques to Fluid Mechanics, Lisbon, Portugal* (Vol. 4).
- Godbersen, P., Gesemann, S., Schanz, D., & Schröder, A. (2024). FlowFit3: Efficient data assimilation of LPT measurements. In *Proceedings of the 21th International Symposium on the Application of Laser and Imaging Techniques to Fluid Mechanics*.
- Harms, T., Brunton, S., & McKeon, B. (2023). Velocity- and derivative-free computation of velocity gradients from Lagrangian particle trajectories. *Proceedings of the 15th International Symposium on Particle Image Velocimetry, ISPIV 2023*.
- Harms, T. D., Brunton, S. L., & McKeon, B. J. (2024). Lagrangian gradient regression for the detection of coherent structures from sparse trajectory data. *Royal Society Open Science*, 11(10).
- Jeon, Y. J., Müller, M., & Michaelis, D. (2022). Fine scale reconstruction (VIC#) by implementing additional constraints and coarse-grid approximation into VIC+. *Experiments in Fluids*, 63(4), 70.
- Kolmogorov, A. N. (1941a). Decay of isotropic turbulence in incompressible viscous fluids. *Dok. Akad. Nauk. S. S. S. R.*, 31, 538-541.
- Kolmogorov, A. N. (1941b). Dissipation of energy in the locally isotropic turbulence. *Dok. Akad. Nauk. S. S. S. R.*, 32, 19-21.
- Kolmogorov, A. N. (1941c). The local structure in incompressible viscous fluid for very large Reynolds numbers. *Dok. Akad. Nauk. S. S. S. R.*, 30, 299-303.
- Kolmogorov, A. N. (1962). A refinement of previous hypotheses concerning the local structure of turbulence in a viscous fluid at high Reynolds number. *Journal of Fluid Mechanics*, 13, 82-85.
- Liszka, T., & Orkisz, J. (1980). The finite difference method at arbitrary irregular grids and its application in applied mechanics. *Computers & Structures*, 11(1-2), 83-95.
- Liu, G.-R., & Liu, M. B. (2003). *Smoothed particle hydrodynamics: a meshfree particle method*. World scientific.
- Mongillo, M., et al. (2011). Choosing basis functions and shape parameters for radial basis function methods. *SIAM undergraduate research online*, 4(190-209), 2-6.
- Pirozzoli, S., Modesti, D., Orlandi, P., & Grasso, F. (2018). Turbulence and secondary motions in square duct flow. *Journal of Fluid Mechanics*, 840, 631-655.

- Schneiders, J. F., & Scarano, F. (2016). Dense velocity reconstruction from tomographic PTV with material derivatives. *Experiments in fluids*, 57(9), 139.
- Schröder, A., & Schanz, D. (2023). 3D Lagrangian particle tracking in fluid mechanics. *Annual Review of Fluid Mechanics*, 55(1), 511–540.
- Suchde, P. (2018). *Conservation and accuracy in meshfree generalized finite difference methods*. Fraunhofer Verlag.
- Suchde, P., Leithäuser, C., Kuhnert, J., & Bordas, S. P. (2025). Volume and mass conservation in Lagrangian meshfree methods. *International Journal for Numerical Methods in Engineering*, 126(2), e7657.
- Taylor, G. I. (1935). Statistical theory of turbulence: Parts I-III. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 151.
- Zhou, K., Grauer, S., Schanz, D., Godbersen, P., Schröder, A., Rockstroh, T., ... Wieneke, B. (2024). Benchmarking data assimilation algorithms for 3D Lagrangian particle tracking. In *21st international symposium on the application of laser and imaging techniques to fluid mechanics, lisbon, portugal* (pp. 1–22).
- Zhou, K., & Grauer, S. J. (2023). Flow reconstruction and particle characterization from inertial Lagrangian tracks. *arXiv preprint arXiv:2311.09076*.
- Zhou, K., Li, J., Hong, J., & Grauer, S. J. (2023). Stochastic particle advection velocimetry (SPAV): theory, simulations, and proof-of-concept experiments. *Measurement Science and Technology*, 34(6), 065302.