

Physics-informed neural networks for dense reconstruction of vortex rings from particle tracking velocimetry

Steinfurth, B.; Hassanein, A.; Doan, N. A.K.; Scarano, F.

DOI

[10.1063/5.0212585](https://doi.org/10.1063/5.0212585)

Publication date

2024

Document Version

Final published version

Published in

Physics of Fluids

Citation (APA)

Steinfurth, B., Hassanein, A., Doan, N. A. K., & Scarano, F. (2024). Physics-informed neural networks for dense reconstruction of vortex rings from particle tracking velocimetry. *Physics of Fluids*, 36(9), Article 095110. <https://doi.org/10.1063/5.0212585>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

RESEARCH ARTICLE | SEPTEMBER 05 2024

Physics-informed neural networks for dense reconstruction of vortex rings from particle tracking velocimetry

B. Steinfurth ; A. Hassanein ; N. A. K. Doan ; F. Scarano 



Physics of Fluids 36, 095110 (2024)

<https://doi.org/10.1063/5.0212585>



Physics of Fluids

Special Topic:

John Michael Dealy (1937-2024): Celebrating His Life

Guest Editors: Alan Jeffrey Giacomini and Savvas G. Hatzikiriakos

[Submit Today!](#)

Physics-informed neural networks for dense reconstruction of vortex rings from particle tracking velocimetry

Cite as: Phys. Fluids **36**, 095110 (2024); doi: [10.1063/5.0212585](https://doi.org/10.1063/5.0212585)

Submitted: 5 April 2024 · Accepted: 12 August 2024 ·

Published Online: 5 September 2024



View Online



Export Citation



CrossMark

B. Steinfurth,^{1,a)} A. Hassanein,² N. A. K. Doan,² and F. Scarano²

AFFILIATIONS

¹TU Berlin, Institute of Aeronautics and Astronautics, Marchstr. 12-14, 10587 Berlin, Germany

²TU Delft, Faculty of Aerospace Engineering, Kluyverweg 1, 2629 HS Delft, The Netherlands

^{a)}Author to whom correspondence should be addressed: ben.steinfurth@tu-berlin.de

ABSTRACT

Phase-resolved volumetric velocity measurements of a pulsed jet are conducted by means of three-dimensional particle tracking velocimetry (PTV). The resulting scattered and relatively sparse data are densely reconstructed by adopting physics-informed neural networks (PINNs), here regularized by the Navier–Stokes equations. It is shown that the assimilation remains robust even at low particle densities ($ppp < 10^{-3}$) where the mean particle distance is larger than 10% of the outlet diameter. This is achieved by enforcing compliance with the governing equations, thereby leveraging the spatiotemporal evolution of the measured flow field. Thus, the PINN reconstructs unambiguously velocity, vorticity, and pressure fields, enabling a robust identification of vortex structures with a level of detail not attainable with conventional methods (binning) or more advanced data assimilation techniques (vortex-in-cell). The results of this article suggest that the PINN methodology is inherently suited to the assimilation of PTV data, in particular under conditions of severe data sparsity encountered in experiments with limited control of the seeding concentration and/or distribution.

Published under an exclusive license by AIP Publishing. <https://doi.org/10.1063/5.0212585>

I. INTRODUCTION

The projection of particle tracking velocimetry (PTV) data onto a structured grid (Cartesian grid reduction, CGR) and the recovery of information in-between the measured scattered vectors persist as a challenge ever since the early PTV experiments.¹ The most straightforward method is to collect velocity vectors into bins that yield the local ensemble-average.^{2,3} Alternatively, linear interpolation circumvents the need to select a bin size. These techniques offer the advantage of simplicity and generality of application to data from a broad range of problems. As such, they do not introduce explicit modeling of the underlying physical process.

In the past decade, methods that impose specific constraints based on the fluid dynamic governing equations have emerged. For instance, mass conservation can be imposed by vanishing velocity divergence during CGR.^{4,5} In addition, the momentum equation can be used for the assimilation of particle trajectories obtained from PTV measurements. When rewritten in terms of the vorticity dynamic equation, this approach has led to the Vortex-in-Cell data assimilation method (VIC+,^{6,7} VIC#,⁸ and VIC-TSA⁹). Similarly, penalization of divergence as well as of a residual pertaining to the momentum

equation have led to the FlowFit method, introduced in Ref. 10 and further developed in Ref. 11. These methods succeed in the dense reconstruction of velocity and vorticity fields from experiments performed at various levels of the seeding density. As is the case with techniques that rely on radial basis functions,^{12,13} they provide continuous or functional representations of the flow field, yielding significant advantages when computing derivatives. On the other hand, the integration of boundary conditions both at the edges of the domain (open boundaries) and at fluid–solid boundaries is regarded as challenging.

Alternative to the above-mentioned methods of PTV data assimilation is the use of artificial neural networks, first attempted by Labonté to track particles across two image frames, thus obtaining the underlying velocity field.¹⁴ The approach was applied to simulated images with simplified motion field, showing robustness in the interpolation between particles and yielding low noise levels. More recently, the suitability of a specific network class, namely physics-informed neural networks (PINNs), for the handling of sparse data has been recognized.¹⁵ The PINN methodology has been introduced by Raissi *et al.*^{16,17} who established several relevant aspects, in particular the incorporation of prior system knowledge by enforcing consistency

with the underlying physics. Examples of this method in the field of fluid mechanics are provided in a recent review,¹⁸ including successful solutions of inverse problems where unknown, or hidden, flow properties are assimilated.

PINNs have also been employed to infer velocity and pressure fields from numerical and experimental PTV data. Wang *et al.* examined the PINN assimilation of scattered data obtained from the analytical solution of two-dimensional Taylor's decaying vortices¹⁹ and the direct numerical simulation of a channel flow, respectively.²⁰ Along with synthetic datasets, Clark di Leoni *et al.*²¹ assessed an experimental PTV dataset of a backward-facing step configuration,²² comparing PINN predictions with results of the Constrained Cost Minimization Method.²³ Most recently, Cai *et al.* demonstrated the capability of PINNs to handle particularly sparse velocity fields by providing only a fraction of the vectors measured in a turbulent jet flow.²⁴

While the aforementioned three studies suggest that PINN frameworks are suited to the assimilation of PTV data, the benchmarking is still ongoing considering the relatively recent introduction of this approach. Therefore, emphasis in the present article is placed on the comparison of PINNs with commonly used alternative data assimilation techniques, namely binning and VIC-TSA. The dataset is provided by 3D-PTV measurements of a pulsed jet ($Re \approx 3000$) featuring a variable concentration of flow tracers, formation of ghost particles and image corruption due to laser light reflections, often reported in real-world experiments as opposed to numerical simulations thereof. Reference velocity fields are obtained by high-resolution planar measurements. Two cases are considered to assess the performance for different degrees of data sparsity: first, the phase-locked nature of the experiment is exploited and velocity vectors are accumulated from multiple snapshots, resulting in a mean particle distance on the order of 10% of the outlet diameter ($d_p/D \approx 0.1$); second, individual snapshots are considered, yielding a number of velocity vectors that is two orders of magnitude lower than in the first case (up to $d_p/D \approx 0.4$).

The experimental apparatus and measurement procedure are introduced first, followed by an overview of the assimilation methods compared in this paper (Sec. III). Then, in Sec. IV, design considerations for the PINN are discussed in detail before a comparison between PINN, traditional data binning, and data assimilation with the vortex-in-cell technique is presented in Sec. V. The main learnings and conclusions are provided in Sec. VI.

II. LAYOUT OF EXPERIMENTS

A. Pulsed jet facility and operation parameters

The flow under consideration is a starting, or pulsed, circular jet consisting of air that is forced from an initial state of rest, leading to the generation of vortex rings. A detailed description of the pulsed jet actuator (Fig. 1) is given in previous studies.^{25–27}

A magnetic valve periodically interrupts the supplied mass-flow, allowing defined amounts of air to enter the nozzle before being ejected through the circular orifice. The divergent-convergent nozzle features an inlet of $d_1 = 3.8$ mm, a maximum cross section $d_{\max} = 10$ mm, converging to an outlet diameter of $D = 5$ mm. The cross section is constant over the final 10 mm upstream of the outlet. Under some assumptions,²⁸ a nominal bulk jet velocity in the exit plane can be estimated based on the (constant) supply mass flow, the outlet diameter, and the relative duration where the valve is open. The jet was operated at a bulk velocity $u_{\text{jet}} \approx 10$ m/s, translating into a




FIG. 1. Pulsed-jet actuator used to produce axisymmetric vortex rings that are evaluated in a $10 \times 10 \times 10$ mm³ domain spanning the jet near-field.

Reynolds number of $Re = u_{\text{jet}}D/\nu \approx 3300$. Considering the concept of a formation number,²⁹ maximum-circulation vortex rings can be expected at $L/D \approx 4$. In the present study, $L = tu_{\text{jet}}$ is interpreted as the time-dependent length of a virtual fluid column emerging from the nozzle exit at bulk velocity. This suggests that the vortex ring formation is completed at $t \approx 2$ ms. In order to produce fully developed vortex rings, the pulse duration is set to $t_p = 5$ ms and the delay between pulses is set to $t_{\text{off}} = 15$ ms so as to reduce the interaction between successively generated vortex rings.

B. Reference planar PIV measurements

As a benchmark for the reconstruction methods assessed in this article, high-resolution planar PIV measurements are performed by seeding the air jet with micrometer-sized DEHS tracer particles³⁰ supplied to the compressed air line feeding the jet. The tracer concentration is set by adjusting the flow through the Laskin nozzle-type particle generator. In addition, an unseeded air feed is connected to the supply for a controlled dilution of the seeded air. The jet exhausts inside a confined box with acrylic glass walls of dimensions 500 mm (axial) \times 250 mm (radial). Before initiating the data acquisition, this cavity is filled with a sufficient amount of tracers such that the velocity field in the ambience can be determined in addition to the seeded jet flow. Variations of the number of particles between individual runs occurred due to the sensitivity of the particle concentration toward the duration of seeding injection.

The particles were illuminated inside the jet symmetry plane with an EverGreen200 dual-cavity Nd:YAG laser ($\lambda = 532$ nm, 2×200 mJ pulse energy). A light sheet with a maximum thickness of 0.1D ($\approx 500 \mu\text{m}$) was formed by means of beam-expanding optics and a knife edge filter. The light scattered by the particle tracers was recorded in double-frame mode with an sCMOS camera (2560×2160 px, 6.5 mm pixel pitch) equipped with an $f = 105$ mm Canon objective and an f-stop set to 16. The images were pre-processed by background removal, subtracting the minimum intensity at each pixel. The field of view of (30×26) mm² (ca. $6D \times 5D$) covers the jet near field, where the vortex ring formation occurs. Each measurement comprises 100 recordings at a selected phase as set by synchronizing the PIV

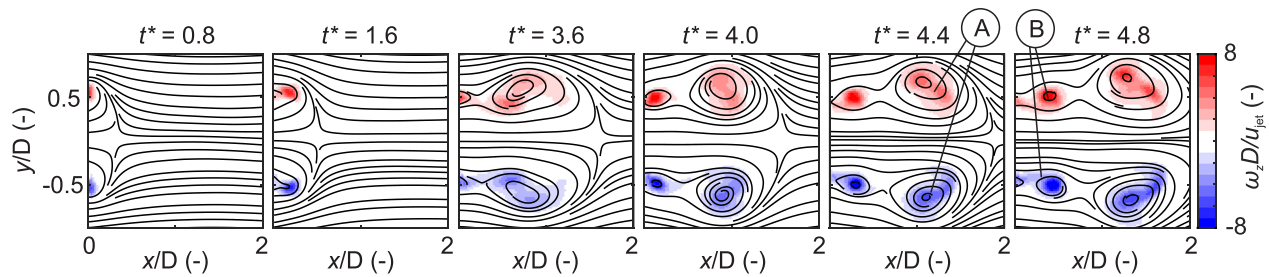


FIG. 2. Vortex ring formation: out-of-plane vorticity overlaid with streamlines in vortex ring frame of reference.

acquisition to the signal controlling the magnetic valve. The time separation of $12\ \mu\text{s}$ corresponds to a particle displacement of $0.12\ \text{mm}$ (9.6px) at the bulk jet velocity. The particle motion was obtained using multi-pass cross correlation to allow the accurate reconstruction of the velocity gradient in the vortex core and across the shear layer. The interrogation window is refined from $96 \times 96\ \text{px}^2$ down to $16 \times 16\ \text{px}^2$, the latter of which corresponds to a spatial resolution of $\Delta x = \Delta y \approx 0.2\ \text{mm}$. An overlap factor of 75% reduces the vector pitch to approximately $0.05\ \text{mm}$ ($0.01D$).

The dynamical evolution of the jet and its vortex ring is illustrated in Fig. 2. The appearance and growth of the axisymmetric vortex is highlighted with the out-of-plane (viz. azimuthal) vorticity. Note that the non-dimensional formation time $t^* = u_{\text{jet}} t / D$ that is referred to in Fig. 2 and throughout this article is defined with reference to the nominal jet velocity and the time coordinate, where $t = 0\text{s}$ marks the moment when the jet emerges from the exit plane ($x = 0\text{mm}$). To ease the topological analysis, streamlines are overlaid after Galilean transformation to a frame of reference that moves with the average vortex ring velocity. With such choice, the saddle point is visualized that separates the rotational fluid issued by the jet on the one hand and the ambient fluid on the other.

At $t^* = 1.6$, a first vortex core (A) detaches from the outlet, followed by a second one (B) at $t^* = 4$.

C. 3D particle tracking

Volumetric velocity measurements of the pulsed jets are performed using a particle tracking velocimetry technique. The experimental setup is schematically represented in Fig. 3.

The seeding procedure is the same as that for the planar PIV measurements; however, the tracer concentration was set to a lower value to mitigate the occurrence of ghost particles. This was achieved by carefully regulating the seeding generator and switching it off just prior to data acquisition. As a result, the seeding concentration features a temporal decay followed by a period when it remained at a relatively stable level for a duration that was sufficiently long to record the data for one time step/phase.

The laser beam of $7\ \text{mm}$ diameter was expanded into an elliptical cross section of approximately $40 \times 20\ \text{mm}$ in the axial and radial (transverse) direction, respectively, which exceeded the dimensions of the region of interest. The Gaussian light distribution is sharply cut by a knife edge filter yielding uniform illumination throughout the measurement domain (Fig. 3, upper left), facilitating particle detection. The light intensity is amplified with a double-pass system,³¹ composed of a planar surface-coated mirror at the outer edge of the confinement

box, reflecting the collimated light back through the region of interest. This approach produces the additional advantage that two directions of forward-scattering are obtained for the benefit of the collected intensity of the imagers positioned along the arc as shown in Fig. 3. The resulting illuminated domain extends over a region of $(30 \times 25 \times 10)\text{mm}^3$ or $(6 \times 5 \times 2)D^3$.

A set of four imagers is placed subtending a tomographic aperture of 50° . The f-stop of the $f = 105\ \text{mm}$ Nikon objectives was set to 32 for all cameras, ensuring a depth of focus encompassing the intersection of the cameras lines of sight with the illuminated region. The 3D system calibration was based on a pinhole model and the coefficients were obtained using a two-level calibration target (LaVision type 11 plate). The volumetric self-calibration procedure³² reduces residual calibration errors below 0.1px .

The temporal (viz. phase) evolution of the jet and the vortex ring formation were resolved by measurements at 31 phases, corresponding to time increments of $100\ \mu\text{s}$ ($\Delta t^* = 0.2$) in the range $t^* = [0, 4]$ and $200\ \mu\text{s}$ ($\Delta t^* = 0.4$) in the range $t^* = [4, 8]$. While approximately constant for each phase (measurement set), the seeding density varied in the range $\text{ppp} = 10^{-4}, \dots, 10^{-3}$ across different sets. Representative raw image samples are shown in Fig. 4 (top) along with the results of pre-processing operations including a minimum-intensity subtraction and setting all intensities below a certain threshold to zero.

Double images were separated by a time delay of $20\ \mu\text{s}$ for PTV data acquisition. The 3D particle detection and motion analysis was performed using the two-pulse variant³³ of the shake-the-box (STB)




FIG. 3. PTV setup; detailed view in the top left shows how the original light sheet highlighted by the green ellipse is cropped.




FIG. 4. Representative particle images for three phases (inverted)—top: raw images; bottom: pre-processed images.

algorithm available in the LaVision *DaVis 10* software. The processing domain was restricted to the jet near field spanning $(20 \times 10 \times 10) \text{ mm}^3$ $[(4 \times 2 \times 2)D^3]$. Particles were detected by limiting triangulation errors to 1px and applying three iterations of particle reconstruction and tracking.³⁴ Particle track detection was restricted to the velocity ranges $u = [-5, 15]$ and $[v, w] = [-5, 5]$ m/s, conservatively spanning the expected velocities for all phases of the jet pulsation.

However, this setting increases the probability of spurious particle pairings and produces some outliers that are removed using a spatial median filter.³⁵

Example particle distributions obtained at three different seeding conditions (the same as in Fig. 4) are shown in Fig. 5 (color-coded by the axial velocity). Outliers, mainly characterized by a streamwise velocity exceeding u_{jet} , are included in the top row whereas the images

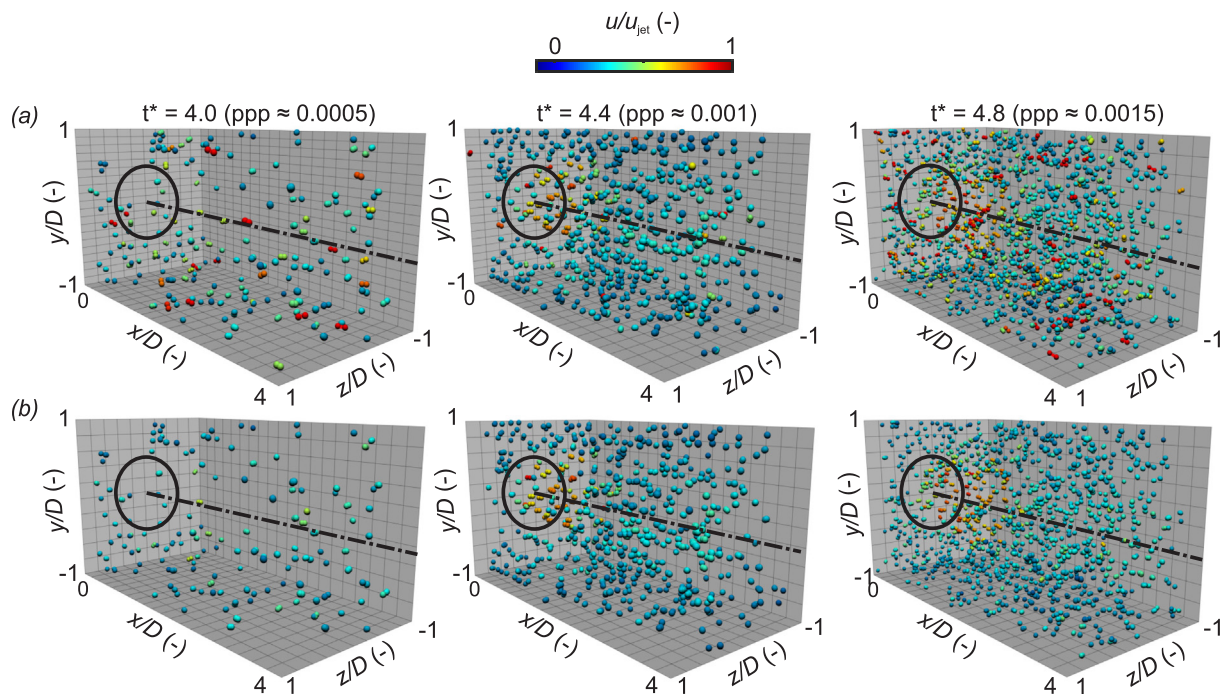


FIG. 5. Examples of particle pairs obtained at different phases and seeding levels (left to right). Top row, includes outliers; bottom row, after data validation; jet outlet and axis indicated by ellipse and dash-dotted line, respectively.

in the bottom row show the filtered PTV results that are used for the data assimilation. The measurement captures the maximum velocity near the jet centerline ($z = y = 0D$), whereas the outer jet region features smaller displacements. The variation of seeding concentration is illustrated in this example. The maximum concentration $\text{ppp} \approx 10^{-3}$ (right-hand side in Fig. 5) is still far below the concentration of $\text{ppp} = 0.1$ that can be handled by the STB algorithm.^{36,37} It should, however, be retained in mind that the latter limit applies to optimal seeding and imaging conditions. In the present experiments, for instance, the particle image diameter is approximately 6.5 pixels, which lowers the upper limit for the concentration.

A priori and *a posteriori* uncertainty analyses for the PTV measurements described above are reported in Appendix B.

D. Datasets for velocity reconstruction

Emphasis is put on the reconstruction capability for different degrees of data sparsity. Therefore, two types of datasets are assessed in this article.

For the statistical dataset (addressed in Sec. V A), the effective spatial concentration of the velocity field is increased by accumulating particle tracks from 100 snapshots taken at the same phase of the pulsed jet. The approach is justified considering that the cycle-to-cycle velocity fluctuations do not exceed 5% of the jet exit velocity (see Appendix A). However, these variations introduce a non-negligible divergence in the accumulated statistical dataset which poses challenges to the reconstruction techniques.

For a single-snapshot dataset (Sec. V B), the number of velocity vectors is two orders of magnitude lower since only individual snapshots are evaluated.

The average number of detected particle tracks at each phase is shown in Fig. 6, yielding significant variations due to slightly different durations for which the measurement domain was seeded prior to acquisition. Since the data for each phase were recorded with one run (10 s), the per-phase fluctuation is relatively low as indicated by the bars in Fig. 6. For both datasets, three phases are selected for further analysis (highlighted by red circles) spanning the range of detected particles $N_p = [100, 600]$, corresponding to the time steps presented in Figs. 4 and 5. Along with the number of particles projected onto the image sensor, ppp , the mean particle distance $d_p = \sqrt[3]{V/N_p}$ is stated in Fig. 5 and indeed throughout the article. Note that $V = 1000 \text{ mm}^3$ is the volume of the domain of interest (see Fig. 1) evaluated with data assimilation techniques.

In addition to the temporal variation of the tracer concentration, the particle distribution was also slightly inhomogeneous (Fig. 7). This




FIG. 6. Average number of particle tracks, particle image density, and mean particle distance per snapshot; error bars span one standard deviation; red circles highlight the conditions considered in the present article.




FIG. 7. Time-averaged mean particle distance at different phases—top row: jet symmetry plane ($z = 0D$); bottom row: cross section at $x = D$.

is reflected in a larger time-averaged particle distance near the jet centerline.

The conditions described above, including low particle concentration and inhomogeneous distribution, are found in many real-world experiments where they typically cannot be controlled to perfection.

III. DENSE VELOCITY RECONSTRUCTION

Three methods are considered to process the 3D-PTV data in the jet near field that extends over a region of $(2 \times 2 \times 2)D^3$ or $(10 \times 10 \times 10)\text{mm}^3$ indicated in Fig. 1.

It is important to note that data binning is applied to velocity data pertaining to single phases whereas short time series spanning three phases of the jet development process are provided for PINN training and vortex-in-cell assimilation.

A. Data binning

Partitioning the domain into sub regions (bins) where the velocity vectors are ensemble-averaged is among the simplest approaches to CGR. Estimating the velocity spatial distribution within the bin by a polynomial function (typically linear or quadratic) fitting the velocity samples reduces the effect of spatial averaging and produces more accurate estimates of the turbulence statistics.³ The resulting spatial resolution depends primarily upon the bin size while the accuracy depends on the tracking precision, the local level of fluctuations, and the number of samples captured in the bin. Consequently, a higher seeding concentration allows to choose smaller bin dimensions.

In the present study, cubic bins with edge lengths of 2 mm (statistical dataset) and 4 mm (single-snapshot dataset) were chosen at an overlap of 75%, and a second-order polynomial was used as fitting

function inside the bins. As a result, the data are represented on a grids of $21 \times 21 \times 21$ and $11 \times 11 \times 11$ points spaced by 0.5 and 1 mm, respectively. The runtime required to process one snapshot is less than a second using an Intel(R) Xeon(R) CPU at a clock speed of 2.7 GHz.

B. Vortex-in-cell assimilation

More advanced CGR methods follow the vortex-in-cell (VIC) paradigm³⁸ by considering the governing equations in vorticity-velocity formulation. While initially used to enhance the temporal resolution of tomographic PIV data,⁶ the method was proven to be applicable to leverage time-resolved recordings for spatial interpolation of scattered data.⁷ The VIC+ algorithm assimilates the instantaneous particle velocity and acceleration taking into account the vorticity transport equations. The method was shown to significantly improve upon tomographic PIV and interpolators to the point of obtaining estimates of the turbulent dissipation rate with reasonable accuracy.⁷ The vortex-in-cell time-segment assimilation (VIC-TSA) leverages the temporal evolution by taking into account time-resolved PTV data to assimilate instantaneous velocity fields.⁹ Its main drawback lies in a sensitivity toward data at the boundaries of the domain, which was recently addressed by introducing some ingenious modifications (VIC#).⁸

In the present study, the VIC-TSA technique is employed to reconstruct the velocity field on $17 \times 19 \times 17$ (statistical dataset) and $9 \times 10 \times 9$ (single-snapshot dataset) grids by applying 40 iterations per snapshot. Using an Intel(R) Xeon(R) CPU at a clock speed of 2.7 GHz, the runtime required for single snapshots is on the order of a few seconds.

C. PINN

The proposed PINN approach features multilayer perceptrons that are trained to model the function

$$\Psi(\tilde{x}, \tilde{y}, \tilde{z}, \tilde{t}) = (\tilde{u}, \tilde{v}, \tilde{w}, \tilde{p}), \quad (1)$$

mapping the input layer consisting of Cartesian coordinates x, y, z and time t to the three velocity components u, v, w and pressure p (Fig. 8). Note that a tilde indicates normalization to the range $[0, 1]$ to avoid




FIG. 8. Schematic representation of PINN—orange: input coordinates (that can be chosen arbitrarily); purple: hidden layers; red: output quantities; green: automatic differentiation.

training issues that could arise from network input/output differing by orders of magnitude.

A multilayer perceptron, also called feedforward neural network,³⁹ consists in a series of layers of neurons connected to all neurons of the subsequent layer through a weight matrix \mathbf{W}_l . The intermediate output of layer l is then expressed as $\xi_l = \mathbf{W}_l^T \xi_{l-1} + \mathbf{b}_l$, where ξ_{l-1} is the output of the previous layer and \mathbf{b}_l is the bias in layer l .⁴⁰ Finally, non-linear behavior is introduced through the activation g , as $\xi_l = g(\xi_l)$. The weights and biases of all hidden layers are updated during training to minimize a loss function. In the case of PINNs, the latter also accounts for physical information.

In the present study, the loss function consists of three parts: \mathcal{L}_{PTV} which accounts for the training data (the PTV data here), \mathcal{L}_{NS} which represents the residual of the governing equations, and \mathcal{L}_{BC} which accounts for the boundary conditions:

$$\mathcal{L} = \mathcal{L}_{\text{PTV}} + \lambda_{\text{NS}} \mathcal{L}_{\text{NS}} + \mathcal{L}_{\text{BC}}. \quad (2)$$

Only one weighting coefficient λ_{NS} is applied in the loss function to balance the physics-informed loss as will be explained in Sec. IV. No further weight is required since the minimization of \mathcal{L}_{BC} is not expected to interfere with the remaining loss terms.

The term \mathcal{L}_{PTV} is the *prediction error* reflecting the deviation between the PINN output and the available experimental data that are represented by velocity vectors $\mathbf{u}_{\text{PTV}} = (u_{\text{PTV}}, v_{\text{PTV}}, w_{\text{PTV}})$ measured at locations $\mathbf{x} = (x, y, z)$. After data normalization, the PTV loss is defined in a mean-squared sense as

$$\mathcal{L}_{\text{PTV}} = \frac{1}{N_p} \sum_{i=1}^{N_p} (\tilde{\mathbf{u}}(\mathbf{x}_i, t_i) - \tilde{\mathbf{u}}_{\text{PTV}}(\mathbf{x}_i, t_i))^2. \quad (3)$$

In the above equation, $(\cdot)^2$ is an inner product and $\tilde{\mathbf{u}}(\mathbf{x}_i, t_i)$ is the output of the network estimated at N_p space-time locations where PTV observational data $\tilde{\mathbf{u}}_{\text{PTV}}(\mathbf{x}_i, t_i)$ are available. The prediction error \mathcal{L}_{PTV} , therefore, allows the network to *anchor* its output at the measurement locations.

The second term, \mathcal{L}_{NS} , accounts for the agreement of the PINN output with the incompressible Navier–Stokes equations evaluated at a set of N_c randomly distributed space-time locations (called collocation points). It is estimated as

$$\mathcal{L}_{\text{NS}} = \frac{1}{N_c} \sum_{j=1}^{N_c} \mathcal{R}^2(\tilde{\mathbf{u}}(x_j, t_j), p(x_j, t_j)), \quad (4)$$

where \mathcal{R}^2 denotes the sum of the squared residuals related to mass conservation and the momentum equations, respectively (after denormalization):

$$\mathcal{R}^2(\mathbf{u}, p) = \|\nabla \mathbf{u} + \partial_t \mathbf{u}\|^2 + \|\mathbf{u} \nabla \mathbf{u} + (\nabla p)/\varrho - \nu \nabla^2 \mathbf{u}\|^2. \quad (5)$$

In the above equation, the required gradients of \mathbf{u} and p with respect to space and time are readily obtained using automatic differentiation.⁴¹ The air density and kinematic viscosity are set to $\varrho = 1.25 \text{ kg/m}^3$ and $\nu = 1.5 \times 10^{-5} \text{ m}^2/\text{s}$, respectively, corresponding to laboratory conditions during experiments. It is worth mentioning that the PINN function [Eq. (1)] can be evaluated at spatiotemporal locations other than those of the training data, and therefore, the collocation points can be arbitrarily sampled so that they ensure that the PINN satisfies the Navier–Stokes equations densely throughout the

space-time domain and not just at the locations where PTV measurements are available. This effectively enables PINNs to increase the resolution of the original (sparse) PTV data by providing a physics-based method to interpolate in-between PTV measurements (in space and time). A parallel can be established between the collocation points and the choice of the output grid in the VIC method.⁷ Furthermore, the space-time coverage of the PINN method can be regarded as an equivalent of the time segment optimization in the VIC-TSA technique⁹ as both methods take into account the temporal evolution of the flow during the measured time interval. While doing so, neither approach relies on Lagrangian information in the shape of particle trajectories.

It is worth mentioning that the pressure (gradient) field is solely assimilated by minimizing the physics-informed residuals. Contrasting with the velocity field, no reference training data are provided. The third term in Eq. (2), \mathcal{L}_{BC} , ensures that the PINN output satisfies boundary conditions including a base value of zero pressure along with vanishing velocity at a transverse jet centerline distance of three outlet diameters. These conditions were enforced at a set of $N_{BC} \approx 4000$ locations surrounding the measurement domain, and the loss term was defined by mean squared deviations ($\tilde{\mathbf{u}}_0$ and \tilde{p}_0 correspond to $\mathbf{u} = 0$ m/s and $p = 0$ Pa when scaled to the range $[0, 1]$):

$$\mathcal{L}_{BC} = \frac{1}{N_{BC}} \sum_{k=1}^{N_{BC}} (\tilde{\mathbf{u}} - \tilde{\mathbf{u}}_0)^2 + \frac{1}{N_{BC}} \sum_{k=1}^{N_{BC}} (\tilde{p} - \tilde{p}_0)^2. \quad (6)$$

The overall loss [Eq. (2)] was minimized using the ADAM optimizer⁴² at a learning rate of $lr = 1 \times 10^{-3}$ for 1000 epochs. Subsequently, to further fine-tune the network, 15 000 evaluations of the loss function were handled by a Limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) optimizer⁴³ driving the supremum norm of the loss function gradient below $\|\nabla(\mathcal{L}_{PTV} + \mathcal{L}_{NS} + \mathcal{L}_{BC})\|_{\infty} = 1 \times 10^{-8}$ for all cases presented in Sec. V. The runtime required for training both on the statistical and the single-snapshot dataset is approximately 15 min using a V100 Tensor Core GPU.

IV. PINN DESIGN CONSIDERATIONS

From the above introduction to PINNs, it is apparent that there are several hyperparameters that can have a large impact on the PINN accuracy in reconstructing the velocity and pressure fields. Specifically, we have observed that the following parameters have a strong influence:

1. Number of collocation points N_c used to estimate the physics-informed loss \mathcal{L}_{NS} [Eq. (4)]
2. The weight of the physics-informed loss λ_{NS} [Eq. (2)]
3. Number of phases provided for training N_t
4. Number of layers/neurons pertaining to the PINN

In Subsections IV A–IV D, the design choices with regard to these parameters are discussed.

A. Number of collocation points

The collocation points play a crucial role in the training of the PINNs as they indicate in which space-time locations the residual of the Navier–Stokes equations is estimated (using the network's prediction). The number of collocation points N_c needs to be chosen sufficiently large to ensure that the PINN output complies with the

governing equations throughout the domain of interest. However, using too many points comes at the cost of unnecessary computational cost.

The sensitivity of the PINN prediction toward the number of collocation points N_c per time step is illustrated in Fig. 9. Here, root-mean-squared deviations between cases with different N_c per training phase N_t are compared to the prediction for the maximum tested $N_c/N_t = 8000$. A discussion on the number of training phases N_t will be provided later (in Sec. III C). Clearly, the output is altered substantially for $N_c/N_t < 1000$ whereas not much is to be gained for larger numbers of collocations points. Therefore, a moderate number of $N_c/N_t = 2000$ is chosen.

B. Weighting the physics-informed loss

The optimization of the PINN output requires attention to balance the loss terms during training by adjusting the weight of the physics-based loss λ_{NS} [Eq. (2)].

In Fig. 10, the output for three settings of the balancing coefficient is presented. This yields three different PINNs, all of which are inferred in the jet symmetry plane at $t^* = 4.0$. For each model, the velocity, vorticity, and the physics-based error fields $\mathcal{R} = \nabla u + \partial_t \mathbf{u} + \mathbf{u} \nabla \mathbf{u} + (\nabla p)/\rho - \nu \nabla^2 \mathbf{u}$ are shown along with the development of loss terms during training (from left to right). The top row represents results for a PINN where the Navier–Stokes equations are not taken into account during training. The result corresponds to an unconstrained interpolation solely driven by the PTV data. The two remaining models were constrained by the Navier–Stokes equations enforced with weights of $\lambda_{NS} = 1 \times 10^{-9}$ (center row) and $\lambda_{NS} = 1 \times 10^{-7}$ (bottom row), respectively.

Recall (from Sec. II B) that at $t^* = 4.0$, a vortex ring is fully developed, pinching off from the trailing jet, featuring a bulk region of large induced velocity enclosed by the toroidal vortex ring corresponding to the maximum vorticity.

The unconstrained solution (top row) yields a consistent velocity field, yet the vorticity pattern is affected by the sparse nature of the measurement. This approach is expected to improve in accuracy when a richer training dataset is provided, in turn requiring a higher seeding concentration. However, the contour plot of the Navier–Stokes residual returns large values in the range $\mathcal{R} = 10^3 - 10^5$, which indicates violations by the interpolated velocity field of the governing Navier–




FIG. 9. Influence of number of collocation points per training phase on velocity predictions.

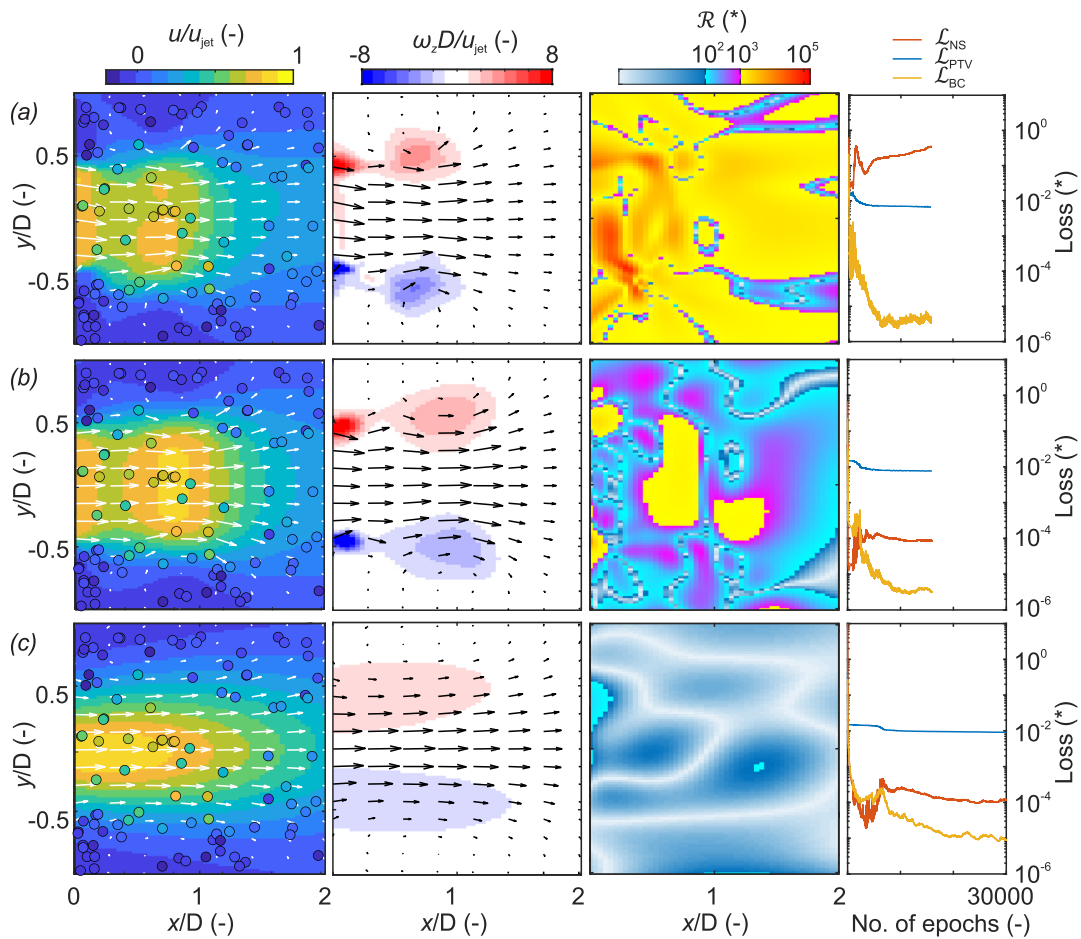


FIG. 10. Influence of physics-informed loss on model output: (a) Navier–Stokes equations not taken into account during training, (b) physics-based loss weighted with $\lambda_{NS} = 1 \times 10^{-9}$, and (c) $\lambda_{NS} = 1 \times 10^{-7}$; the colored circles shown in the left column indicate the PTV measurement data at a maximum distance of 0.1 mm (0.02D) from the symmetry plane.

Stokes equations. Such violations arise not only because the pressure output (not shown here) is completely unconstrained, but also the observational velocity data themselves may be in conflict with the Navier–Stokes equations due to the measurement uncertainty on the one hand and the accumulation of particle tracks from different snapshots on the other. Specifically, the latter can be expected to introduce divergence due to the non-negligible cycle-to-cycle variation. The loss term development during training (right column) shows a monotonic decrease in \mathcal{L}_{PTV} while the physics-based term increases.

In the center row of Fig. 10, a weight of $\lambda_{NS} = 1 \times 10^{-9}$ was chosen for the physics-based loss, leading to smooth velocity and vorticity distributions rather consistent with the reference flow field shown in Fig. 2 (i.e., local extrema associated with the primary and secondary vortex ring cores). The residual also drops below $\mathcal{R} = 10^3$ almost throughout the symmetry plane, and across the entire training domain (not shown here). In accordance, \mathcal{L}_{NS} decreases by approximately four orders of magnitude compared to the unconstrained case.

By further enforcing compliance with the Navier–Stokes equations ($\lambda_{NS} = 1 \times 10^{-7}$, bottom row), the physics-based loss is reduced

to an even greater extent (see second column from the right). However, this condition appears to overly constrain the solution, forcing the PINN output toward a trivial solution of the Navier–Stokes equations (i.e., homogeneous velocity), which greatly departs from the measurement data. In other words, the influence of the PTV data are marginalized by increasing the contribution of the physics-based loss, allowing for larger deviations between the measurement data and the PINN output. This conflict is illustrated in Fig. 11, showing the physics-informed loss (orange curve) and the training data loss (blue curve) for a varied weighting coefficient. By taking the square-root of the training data loss defined in Eq. (3), the values can be interpreted as deviations in (m/s) normalized with the velocity range pertaining to the training dataset.

Considering the pivotal role played by the weight coefficient on the PINN output, one may wonder how to set this parameter sensibly. While acknowledging that this issue is an open research question,^{44–46} a heuristic approach is proposed based upon a crossed evaluation of the PINN output compared to the PTV data (Fig. 10, left column) and the magnitude of the data loss term (Fig. 11). The model output may




FIG. 11. Loss terms as a function of the physics-based weight coefficient; dashed line indicates threshold for PTV data loss informed by measurement uncertainty.

be required to follow the training data within a bound that is defined by the velocity measurement uncertainty. The latter is estimated to be on the order of 10% of the bulk jet velocity by an *a posteriori* analysis (see Appendix B). The largest order of magnitude in λ_{NS} that results in a model with a data loss smaller than 10% is $\lambda_{NS} = 1 \times 10^{-9}$ (center row in Fig. 10). We, therefore, proceed with this value.

C. Number of pulsed jet phases

When designing a PINN for dense flow field reconstruction, a choice needs to be made regarding the length of the modeled time series. For instance, one may use a single PINN for the entire available time series. However, this would require the same set of neurons to represent both the spatial flow geometry and its complex temporal evolution. An alternative is to cut the reconstruction problem into smaller time windows and train a shallower PINN for each window.

The latter approach is taken here as no single PINN could be trained to accurately reconstruct the entire time series. The question is now related to the optimal time window over which to train individual networks. It should be noted first that providing a single time step is not sufficient as the Navier–Stokes equations require a time-derivative information [Eq. (5)] which would not be available in that case. Therefore, PTV data spanning multiple phases need to be provided. The influence of the number of phases on selected model predictions after the same number of training epochs ($N_{LBFGS} = 15\,000$) is shown in Fig. 12 for $N_t = (3, 5, 7)$ (from top to bottom). For each case, the presented target phase is $t^* = 4.0$ but different numbers of preceding and succeeding phases are considered. The number of collocation points is increased by 2000 for each added phase.

Clearly, cases b and c depart from the ground truth, which is manifested in blurred spatial velocity gradients yielding an elongated shear-layer where the primary and secondary vortex rings become

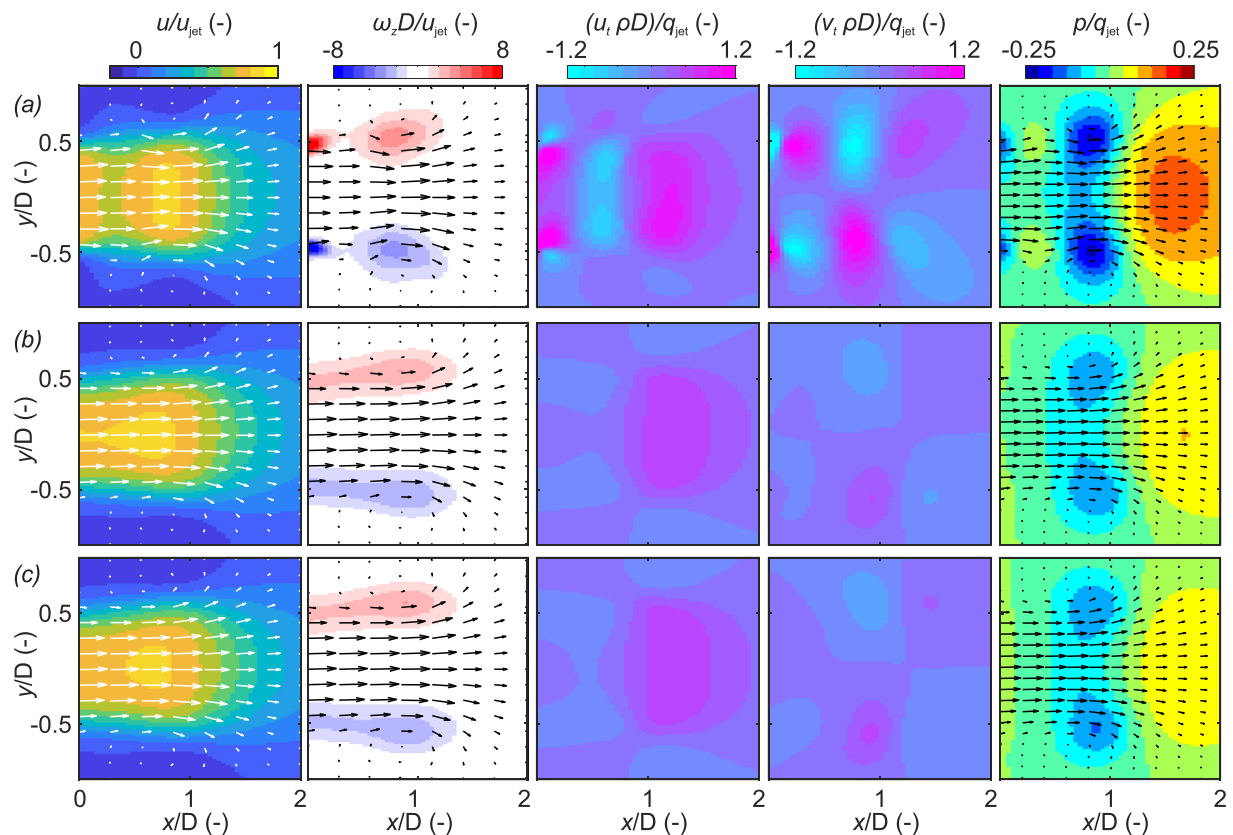


FIG. 12. Influence of number of training phases on velocity, vorticity, and pressure predictions after the same number of epochs: (a) $N_t = 3$, (b) $N_t = 5$, and (c) $N_t = 7$.

indistinguishable. Similarly, the temporal velocity gradients u_t and v_t are diminished, resulting in a reduced magnitude observed in the static pressure field (right column). This is explained by the inability of the PINN to capture the non-linear time evolution beyond a certain length. To handle the larger degree of complexity introduced by adding training phases, one may increase the network dimensions and/or train the PINN for a larger number of epochs. In the present study, the number of epochs required to reach convergence (for a fixed model architecture) was approximately proportional to the number of training phases. However, the computational cost is increased disproportionately as collocation points are added for each phase. We, therefore, chose to train the PINNs on $N_t = 3$ phases for all cases presented in the following where the target phase is accompanied by phases just before and after. It is worth mentioning that this does not mean that the temporal gradients are computed based on only three time steps as the collocation points are sampled randomly inside the range spanned by the preceding and succeeding training phases.

D. Network depth

Finally, the width and depth of the PINN need to be chosen. Keeping the number of neurons per layer constant at $N_n = 50$ (each activated with a tanh function), the number of layers was varied and the overall loss evaluated (Fig. 13).

As the network depth is increased, its expressivity is enhanced, leading to a reduction of the overall loss. Beyond a certain network depth, the number of dataset features not yet represented by the model decreases. In other words, choosing an excessive network depth brings no significant benefits but increases the wall time required for training. The current analysis suggests that increasing the depth beyond 12 layers returns negligible benefits.

V. DEMONSTRATION OF PTV DATA ASSIMILATION

In this section, the PINN method is compared to alternative CGR techniques, namely data binning and VIC-TSA. It is important to reiterate that for the VIC-TSA and PINN methods, measurement data from three time steps are taken into account whereas binning is applied per time step.

The performance of the three methods is examined for different degrees of PTV data sparsity. In the first scenario, the number of particles inside the near-outlet domain $x = 0, \dots, 10$ mm is on the order of $\mathcal{O}(N_p) = 10^4$, which was facilitated by accumulating particle tracks




FIG. 13. Overall loss as a function of the number of hidden layers; training with $N_t = 3$ phases and target phase $t^* = 4.0$.

from 100 snapshots. Then, individual snapshots will be evaluated; hence, the number of particles is two orders of magnitude lower. The particle image density for these two conditions ranges between $\mathcal{O}(\text{ppp}) = 10^{-1}$ (accumulated particles, highest concentration) and $\mathcal{O}(\text{ppp}) = 10^{-5}$ (single snapshot, lowest concentration).

A. Training on a statistical dataset

Along with the reference high-resolution planar PIV measurements (right column), the results for the three methods are displayed in Fig. 14. Recall that for the binning method, cubes with an edge length of 2 mm (75% overlap) were chosen, yielding a vector pitch of 0.5 mm. The same spatial resolution is imposed with the VIC-TSA method. Although the trained PINN can be sampled at arbitrary resolution, an equal vector spacing was chosen for ease of comparison.

For the first phase ($t^* = 4.0$, $\text{ppp} \approx 2 \times 10^{-2}$), the binning of approximately 6700 velocity vectors leads to some discontinuities and a bulk velocity that is smaller than observed in the reference planar PIV results. The VIC-TSA method yields a higher degree of similarity with the planar PIV results. The PINN method also captures the main features of the starting jet. Specifically, there is a connected region of increased axial velocity that is enclosed by two counter-rotating projections of the primary vortex ring onto the symmetry plane. Furthermore, the secondary vortex ring that starts to develop at the jet outlet is captured although its velocity distribution in lateral direction differs from the one indicated by planar PIV measurements. For the subsequent time steps, featuring a higher seeding concentration, a clearer picture is revealed by the binning and VIC-TSA methods, and comparable reconstruction results are noted for all three techniques.

The above discussion is extended to the accuracy and resolution of the spatial velocity gradient by observing the reconstructed out-of-plane vorticity (Fig. 15).

As indicated by planar PIV measurements, two local extrema (primary and secondary vortex rings) can be expected on both sides of the jet symmetry plane. While binning and VIC-TSA reconstructions yield blurred and fragmented distributions for the first time step, respectively, the PINN prediction exhibits more similarities with the reference PIV measurements. Specifically, the vorticity magnitude of the secondary vortex ring is captured more accurately. Yet, the intricate structure of the detached vortex ring, which seems to be composed of two merged vortices, is not rendered in full detail.

Assuming axisymmetric vortex rings with no swirl, the circulation, hydrodynamic impulse, and kinetic energy can be computed based on the velocity and vorticity fields presented above as⁴⁷

$$\Gamma = \iint \omega_z dy dx; \quad I = \rho \pi \iint \omega_z r_y^2 dy dx; \quad E = \rho \pi \iint (u^2 + v^2) r_y dy dx, \quad (7)$$

where r_y denotes the radius of the vortex ring. The integration limits are chosen such that only data associated with the vortex ring are considered. The respective regions are defined by two criteria: (1) they enclose the vortex core and (2) they exhibit vorticity above 10% of the maximum value. Using the invariants of motion stated in Eq. (7), the non-dimensional vortex ring energy

$$\alpha = \frac{E}{\Gamma^{3/2} I^{1/2}}, \quad (8)$$




FIG. 14. Velocity field reconstruction with different methods at 0.5 mm grid resolution, reference planar PIV: $\Delta x = \Delta y = 0.1$ mm, vectors indicated by white arrows with 1 mm spacing: (a) $t^* = 4.0$ ($N_p \approx 6700$, $d_p/D \approx 0.11$), (b) $t^* = 4.4$ ($N_p \approx 27300$, $d_p/D \approx 0.07$), and (c) $t^* = 4.8$ ($N_p \approx 59000$, $d_p/D \approx 0.05$).

is obtained for different time steps. Note that this quantity decreases as the vortex ring grows in thickness during its formation (Fig. 16). The limiting value reported by Gharib *et al.*²⁹ is $\alpha = 0.33$, which is highlighted by the dashed horizontal. Indeed, this value is also indicated by planar PIV measurements in the present study at $t^* > 5$. Whereas a deviation is noticed at smaller formation times, the α curve obtained for the PINN output (blue line) almost collapses with the planar PIV data otherwise. In contrast, slightly larger differences are on display for the binning and VIC-TSA methods where the non-dimensional energy does not drop below $\alpha = 0.4$. We, therefore, conclude that the integral vortex ring properties are represented most accurately by PINN models.

It is worth recalling that all methods have been applied to volumetric PTV data but only a slice of the respective output, namely the symmetry plane, has been presented so far. Next, in Fig. 17, the full three-dimensional structure of the flow field is represented by means of iso-surfaces. The Q criterion⁴⁸ (red iso-surface) is adopted to unambiguously detect the vortex ring produced by the pulsed jet. Furthermore, the PINN also returns the spatial distribution of the relative static pressure ($\Delta p = -11$ Pa), illustrated by a blue iso-surface in the last column. The latter is strongly correlated with the vortex core and it has been often considered as equivalent vortex identification

criterion.⁴⁸ Recall that the pressure gradient is assimilated through the momentum equation and the reference value for pressure is defined via boundary conditions three outlet diameters away from the jet exit axis ($y = z = \pm 3D$, $x = [0, 4]D$).

For the PINN, a smooth primary vortex is revealed by the Q criterion, even at $t^* = 4.0$ (i.e., for the condition of highest data sparsity). Furthermore, the formation of the secondary vortex ring can be observed which is not captured by the other methods. At $t^* = (4.4, 4.8)$, both vortex rings are returned by all methods although the lowest noise level appears to be at hand for the PINN method.

In the pressure field predicted by the PINN, toroidal structures corresponding to the primary and secondary vortex rings are adequately captured at $t^* = (4.4, 4.8)$ while an oblate spheroid and a fragmentary ring are produced at the first phase. Against this backdrop, a reasonable assimilation of the pressure field, being a hidden quantity, can be attested.

B. Training on a single snapshot

In the following, even more challenging conditions for data assimilation are considered. Specifically, only vectors from single snapshots, as opposed to the dataset based on 100 snapshots assessed in the previous section, are assimilated. While the evaluated time steps

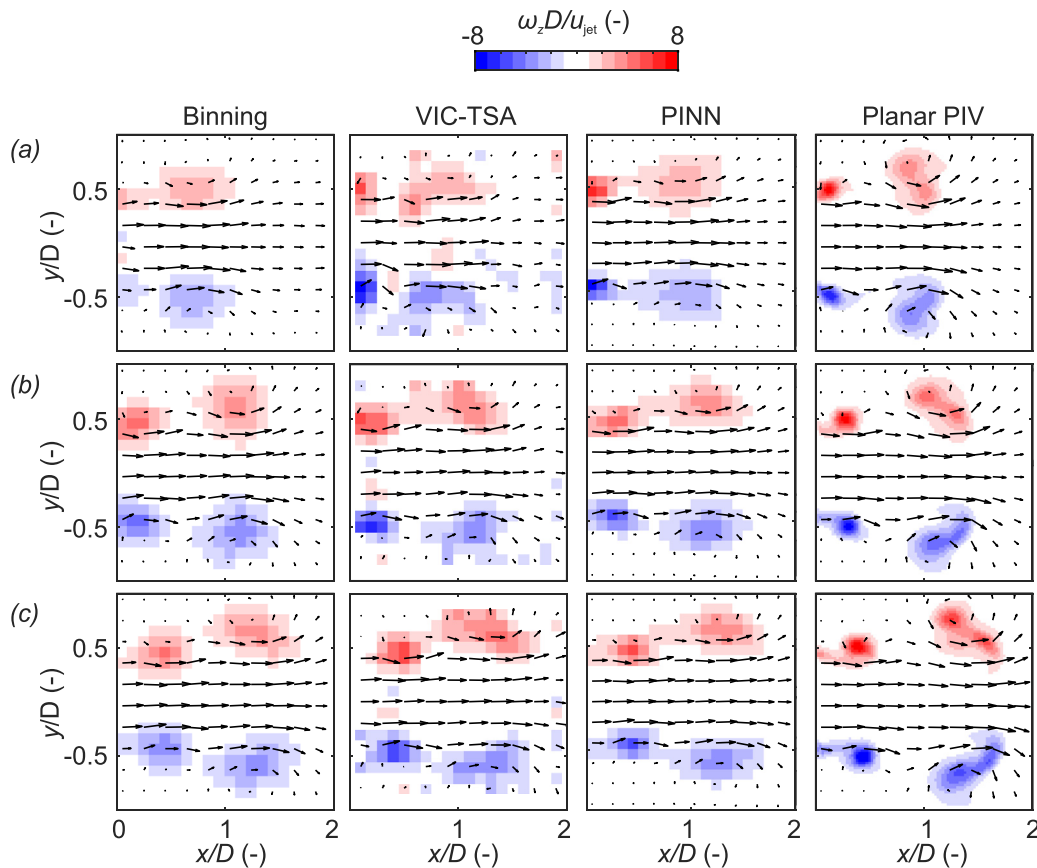


FIG. 15. Vorticity fields obtained by different methods compared to reference planar PIV results, grid resolution, and vector spacing the same as in Fig. 14, (a)–(c) same as in Fig. 14.

remain the same, the number of particles is, therefore, reduced by two orders of magnitude. To cope with these data sparsity, the bin size was increased from 2 to 4 mm leading to a bin volume increase by a factor 8. For ease of comparison, the resolution pertaining to VIC-TSA and




FIG. 16. Development of non-dimensional vortex ring energy obtained by different methods compared to reference planar PIV results.

PINN output are increased by the same factor. The assimilation results for this single-snapshot dataset are presented in Fig. 18.

At the lowest particle concentration ($N_p = 89$, i.e., less than 0.1 particles per mm^3), a substantial performance deficit is observed at the first phase. For the binning method and VIC-TSA, the jet flow is barely recognizable. The PINN, on the other hand, predicts a velocity distribution that is similar to the reference velocity field, apparently benefitting to a greater extent than VIC-TSA from the availability of data at the preceding and succeeding phases ($N_p = 137$ and $N_p = 389$, respectively).

As can be expected, the binning and VIC-TSA methods yield more reliable results for the velocity field at $t^* = (4.4, 4.8)$ than at $t^* = 4.0$ given the larger number of particles for the two later time steps. Yet, the largest similarity with the true velocity field is achieved by the PINN. This also applies to the vorticity fields where the PINN captures both the primary and secondary vortex ring and, perhaps surprisingly, the deviation from the test case with a much larger number of particles (Fig. 15) is relatively small.

Finally, the same type of iso-contours as in Fig. 17 are presented for the evaluation of single snapshots (Fig. 19).

As was indicated by velocity and vorticity fields presented above, binning and VIC-TSA produce more noisy vortex rings at $t^* = (4.4, 4.8)$ whereas no such structure can be found at $t^* = 4.0$. The PINN output allows for a clearer identification of the same flow structures.




FIG. 17. Time series of iso-contours used for vortex detection obtained by different methods: (a) $t^* = 4.0$ ($N_p \approx 6700$, $d_p/D \approx 0.11$), (b) $t^* = 4.4$ ($N_p \approx 27300$, $d_p/D \approx 0.07$), and (c) $t^* = 4.8$ ($N_p \approx 59000$, $d_p/D \approx 0.05$): Q criterion shown at $Q = 5 \cdot 10^6 \text{ s}^{-2}$ (red) and low pressure region at $\Delta p = -11 \text{ Pa}$ (blue).

To summarize this section, we conclude that assimilating PTV data using PINNs appears to be superior to the binning method and data assimilation using VIC-TSA in the present study. Specifically, smaller amounts of velocity data are required to infer reasonable velocity fields of larger spatial coherence.

VI. CONCLUSIONS

The objective of this study was to assess the suitability of PINNs to assimilate measurement data obtained by PTV, an approach recently proposed by Wang *et al.*,²⁰ Clark di Leoni *et al.*,²¹ and Cai *et al.*²⁴

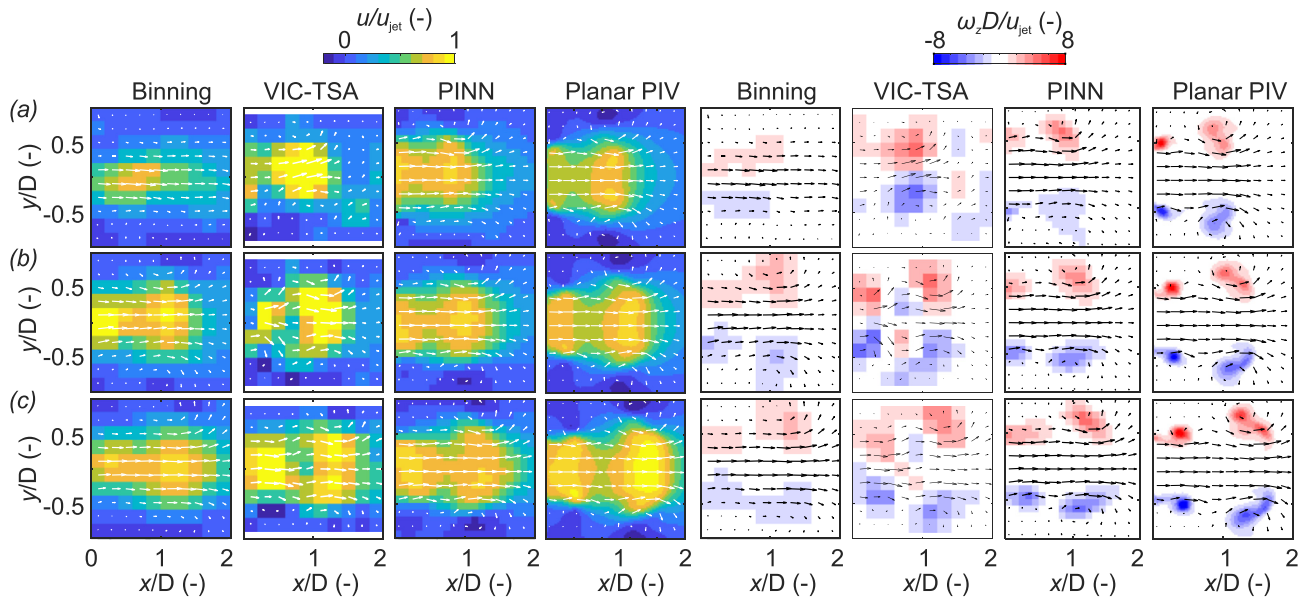


FIG. 18. Velocity and vorticity field reconstruction with different methods at 1mm grid resolution, reference planar PIV: $\Delta x = \Delta y = 0.1 \text{ mm}$, vectors indicated by white arrows with 1mm spacing: (a) $t^* = 4.0$ ($N_p \approx 100$, $d_p/D \approx 0.43$), (b) $t^* = 4.4$ ($N_p \approx 400$, $d_p/D \approx 0.27$), and (c) $t^* = 4.8$ ($N_p \approx 700$, $d_p/D \approx 0.23$).




FIG. 19. The same as in Fig. 17, but now iso-surfaces obtained from a single snapshot: (a) $t^* = 4.0$ ($N_p \approx 100$, $d_p/D \approx 0.43$), (b) $t^* = 4.4$ ($N_p \approx 400$, $d_p/D \approx 0.27$), and (c) $t^* = 4.8$ ($N_p \approx 700$, $d_p/D \approx 0.23$).

PTV delivers scattered velocity field information associated with the displacement of individual particles supplied to the flow. Considering the configuration of pulsed jets, PINNs were set up to predict three-component velocity vectors along with the pressure at (arbitrary) query locations, here defined by Cartesian coordinates and time. This mapping was learned by constraining the PINN output by time series of measurement data on the one hand and the three-dimensional, incompressible Navier–Stokes equations on the other.

It was shown that care must be taken when weighting the major loss terms (i.e., the PTV data loss and the physics-based loss). Specifically, only taking into account the measurement data leads to model predictions violating the governing partial differential equations, for which two sources can be identified: first, a lack of constraints for the interpolation in-between the relatively sparse measurement data, especially in terms of the pressure output; second, the introduction of divergence in the case of the statistical database that arises from accumulating velocity data from different snapshots with slight variations. The PINN paradigm is suited to handle both of these issues. However, overly enforcing physics compliance yields trivial solutions with blurred features of the flow at hand. To handle this balancing act, we propose to evaluate the training data loss for a range of weighting coefficients. As a larger relative importance is assigned to the physics-based loss, the data loss increases, and the order of the velocity measurement uncertainty can be applied as a reasonable threshold for the latter.

One major focus point of this study was the comparison of the PINN method with two alternative standard techniques, namely binning³ and a vortex-in-cell technique (VIC-TSA⁹). In terms of the required computational cost, the binning method is the least expensive

as one snapshot is processed in less than a second using an Intel(R) Xeon(R) CPU at 2.7 GHz clock speed. For the same amount of data, the VIC-TSA processing takes approximately 4 s whereas a PINN model is trained in approximately 15 min (on a V100 Tensor Core GPU).

While both binning and VIC-TSA perform reasonably well in cases of higher particle numbers, they fail to overcome the data sparsity otherwise. The same is not true for the PINN that is shown to adequately replicate the flow pattern underlying the particle-based measurements. Consequently, a clearer picture of the occurring vortex structures is revealed, even for cases of substantial data sparsity. As an example, clear iso-surfaces of the Q criterion corresponding to the studied vortex rings are delivered at a mean particle distance as low as $d_p \approx 0.25D$ ($ppp < 10^{-3}$). In addition, the PINN allows to infer regions of low pressure, assimilated through the Navier–Stokes equations, that can augment the vortex identification. To obtain the pressure magnitude, in addition to the gradient that appears in the Navier–Stokes equations, a base pressure can be seamlessly integrated into the PINN training. As a particular benefit of the PINN approach, this boundary condition can be enforced at locations outside of the measurement domain.

Another advantage of the PINN approach lies in the robust physics-based spatiotemporal interpolation, which is consistent with the concept of “pouring time into space.”⁴⁹ This feature may help alleviate data sparsity (per snapshot) due to low tracer concentrations by providing multiple time steps that capture the temporal evolution. However, in the present study, we noted a performance degradation when the number of time steps was increased beyond a number of three, which was explained by the inability of the network (i.e., a single set of neurons) to capture the dynamics both in space and time. While

noting that up to 1000 time steps were provided successfully in another study,²¹ establishing guidelines for the training of PINNs in multi-scale (in terms of space and time) configurations remains a future research direction.

ACKNOWLEDGMENTS

This work was supported by a postdoc fellowship of the German Academic Exchange Service (DAAD) under Program ID 57664191.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

B. Steinfurth: Conceptualization (equal); Data curation (equal); Funding acquisition (lead); Investigation (equal); Methodology (equal); Software (lead); Validation (equal); Visualization (lead); Writing – original draft (lead); Writing – review & editing (equal). **A. Hassanein:** Investigation (supporting); Methodology (supporting). **N. A. K. Doan:** Conceptualization (equal); Investigation (equal); Methodology (equal); Writing – review & editing (equal). **F. Scarano:** Conceptualization (equal); Investigation (equal); Project administration (lead); Resources (lead); Writing – review & editing (equal).

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

APPENDIX A: REPRODUCIBILITY OF PULSED JET FLOW

As a measure for the reproducibility of the flow, the mean fluctuating part of the main velocity component determined with high-resolution planar PIV is presented in Fig. 20. Recall that for each phase, 100 instantaneous velocity fields are considered.

While velocity fluctuations are small compared to the jet exit velocity (mostly below 3%, i.e., on the order of the planar PIV measurement uncertainty) in the majority of the measurement domain, large fluctuations are observed to encompass the vortex core. However, these fluctuations may also be caused by local maxima of the measurement uncertainty due to a combination of high velocity gradients, vanishing particle displacement and seeding depletion (centrifugal forces) rather than a lack of reproducibility.

Overall, a high degree of reproducibility can be attested, justifying the approach of accumulating particle tracks for the respective phases.

APPENDIX B: UNCERTAINTY ESTIMATION

The uncertainty of velocity information obtained with PTV can be determined *a priori* by propagating the uncertainty associated with the particle location. The latter is determined by $\sigma_x = c_r d_{p,i} / M$, where $c_r = 0.3$ represents the uncertainty in locating the particle centroid,⁵⁰ $d_{p,i} \approx 6.5\text{px}$ is the particle image diameter,

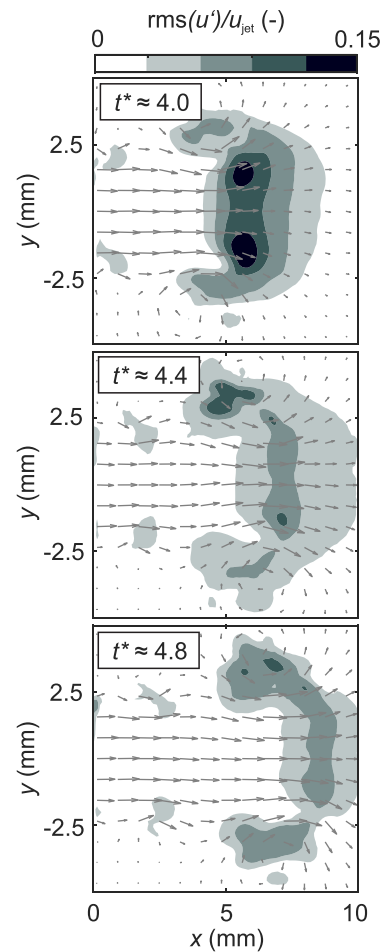


FIG. 20. Time series of pulsed jet—top: phase-averaged axial velocity contour overlaid with in-plane velocity vectors, locations of vortex core highlighted by red and blue circles; bottom: probability density functions of vortex core outlet distances for the displayed phases.

and $M \approx 0.3$ the magnification. This yields $\sigma_x \approx 40\ \mu\text{m}$ such that the uncertainty in velocity measurement becomes $\sigma_u = \sigma_x / \Delta t \approx 2\text{ m/s}$ (i.e., 20% of the bulk jet velocity).

Furthermore, the uncertainty is estimated *a posteriori* by comparing PTV data with a pseudo ground truth (in the absence of an absolute ground truth). The reference velocity field is provided by a well-converged VIC-TSA evaluation applied to a segment of five time steps with $t^* = 4.0$ being the target time step. By providing phase-locked data at the respective phases, this approach also accounts for the uncertainty induced by cycle-to-cycle fluctuations. After linear interpolation at the available particle locations, the disparity is computed, and its dispersion is used to estimate the uncertainty. Specifically, the values corresponding to the half-peak heights are determined for the three velocity components (Fig. 21). The highest values are determined for the u and w components (approximately 1 m/s). The former is driven by the intense axial dynamics of the flow while the latter is explained by the optical setup where the depth position and displacement are determined




FIG. 21. Histograms of disparity between PTV measurement data and pseudo ground truth provided by VIC-TSA for all three velocity components.

less accurately. Furthermore, errors associated with ghost particles are produced mostly along the depth direction, increasing the uncertainty of the w component, in a relative sense. For the v component, the uncertainty is lower in comparison (0.7 m/s).

In summary, the *a posteriori* analysis suggests an uncertainty on the order of 10% of the bulk jet velocity.

REFERENCES

- ¹J. C. Agüi and J. Jiménez, "On the performance of particle tracking," *J. Fluid Mech.* **185**, 447–468 (1987).
- ²N. Kasagi and K. Nishino, "Probing turbulence with three-dimensional particle-tracking velocimetry," *Exp. Therm. Fluid Sci.* **4**, 601–612 (1991).
- ³N. Agüera, G. Cafiero, T. Astarita, and S. Discetti, "Ensemble 3D PTV for high resolution turbulent statistics," *Meas. Sci. Technol.* **27**, 124011 (2016).
- ⁴C. M. de Silva, J. Philip, and I. Marusic, "Minimization of divergence error in volumetric velocity measurements and implications for turbulence statistics," *Exp. Fluids* **54**, 1557 (2013).
- ⁵D. Schiavazzi, F. Coletti, G. Iaccarino, and J. K. Eaton, "A matching pursuit approach to solenoidal filtering of three-dimensional velocity measurements," *J. Comput. Phys.* **263**, 206–221 (2014).
- ⁶J. F. G. Schneiders, R. P. Dwight, and F. Scarano, "Time-supersampling of 3D-PTV measurements with vortex-in-cell simulation," *Exp. Fluids* **55**, 1692 (2014).
- ⁷J. F. G. Schneiders and F. Scarano, "Dense velocity reconstruction from tomographic PTV with material derivatives," *Exp. Fluids* **57**, 139 (2016).
- ⁸Y. J. Jeon, M. Müller, and D. Michaelis, "Fine scale reconstruction (VIC#) by implementing additional constraints and coarse-grid approximation into VIC+," *Exp. Fluids* **63**, 70 (2022).
- ⁹F. Scarano, J. F. G. Schneiders, G. G. Saiz, and A. Sciacchitano, "Dense velocity reconstruction with VIC-based time-segment assimilation," *Exp. Fluids* **63**, 96 (2022).
- ¹⁰S. Gesemann, F. Huhn, D. Schanz, and A. Schröder, "From noisy particle tracks to velocity, acceleration and pressure fields using B-splines and penalties," in 18th International Symposium on Applications of Laser Techniques to Fluid Mechanics, 2016.
- ¹¹F. Ehlers, A. Schröder, and S. Gesemann, "Enforcing temporal consistency in physically constrained flow field reconstruction with FlowFit by use of virtual tracer particles," *Meas. Sci. Technol.* **31**, 094013 (2020).
- ¹²P. Vedula and R. J. Adrian, "Optimal solenoidal interpolation of turbulent vector fields: Application to PTV and super-resolution PIV," *Exp. Fluids* **39**, 213–221 (2005).
- ¹³L. D. C. Casa and P. S. Krueger, "Radial basis function interpolation of unstructured, three-dimensional, volumetric particle tracking velocimetry data," *Meas. Sci. Technol.* **24**, 065304 (2013).
- ¹⁴G. Labonté, "Neural network reconstruction of fluid flows from tracer-particle displacements," *Exp. Fluids* **30**, 399–409 (2001).
- ¹⁵Y. Du, M. Wang, and T. A. Zaki, "State estimation in minimal turbulent channel flow: A comparative study of 4DVar and PINN," *Int. J. Heat Fluid Flow* **99**, 109073 (2023).
- ¹⁶M. Raissi, A. Yazdani, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *Science* **367**, 1026–1030 (2020).
- ¹⁷M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *J. Comput. Phys.* **378**, 686–707 (2019).
- ¹⁸S. Cai, Z. Mao, Z. Wang, M. Yin, and G. E. Karniadakis, "Physics-informed neural networks (PINNs) for fluid mechanics: A review," *Acta Mech. Sin.* **37**, 1727–1738 (2022).
- ¹⁹G. I. Taylor, "On the decay of vortices in a viscous fluid," *London, Edinburgh, Dublin Philos. Mag. J. Sci.* **46**, 671–674 (1923).
- ²⁰H. Wang, Y. Liu, and S. Wang, "Dense velocity reconstruction from particle image velocimetry/particle tracking velocimetry using a physics-informed neural network," *Phys. Fluids* **34**, 017116 (2022).
- ²¹P. Clark Di Leoni, K. Agarwal, T. A. Zaki, C. Meneveau, and J. Katz, "Reconstructing turbulent velocity and pressure fields from under-resolved noisy particle tracks using physics-informed neural networks," *Exp. Fluids* **64**, 95 (2023).
- ²²S. Gopalan and J. Katz, "Flow structure and modeling issues in the closure region of attached cavitation," *Phys. Fluids* **12**, 895 (2000).
- ²³K. Agarwal, O. Ram, J. Wang, Y. Lu, and J. Katz, "Reconstructing velocity and pressure from noisy sparse particle tracks using constrained cost minimization," *Exp. Fluids* **62**, 75 (2021).
- ²⁴S. Cai, C. Gray, and G. E. Karniadakis, "Physics-informed neural networks enhanced particle tracking velocimetry: An example for turbulent jet flow," *IEEE Trans. Instrum. Meas.* **73**, 2519109 (2024).
- ²⁵B. Steinfurth and J. Weiss, "Vortex rings produced by non-parallel planar starting jets," *J. Fluid Mech.* **903**, A16 (2020).
- ²⁶B. Steinfurth and J. Weiss, "Velocity ratio effect on flow structures of non-parallel planar starting jets in cross-flow," *J. Fluid Mech.* **915**, A11 (2021).
- ²⁷B. Steinfurth and J. Weiss, "Modelling the decay of finite-span starting and stopping wall jets in an external stream," *J. Fluid Mech.* **951**, A28 (2022).
- ²⁸B. Steinfurth, "Flow physics of pulsed-jet actuation," Ph.D. thesis (TU Berlin, 2022).
- ²⁹M. Gharib, E. Rambod, and K. Shariff, "A universal time scale for vortex ring formation," *J. Fluid Mech.* **360**, 121–140 (1998).
- ³⁰C. J. Kähler, B. Sammler, and J. Kompenhans, "Generation and control of tracer particles for optical flow investigation in air," *Exp. Fluids* **33**, 736–742 (2002).
- ³¹S. Ghaemi and F. Scarano, "Multi-pass light amplification for tomographic particle image velocimetry applications," *Meas. Sci. Technol.* **21**, 127002 (2010).
- ³²B. Wieneke, "Volume self-calibration for 3D particle image velocimetry," *Exp. Fluids* **45**, 549–556 (2008).

- ³³M. Novara, D. Schanz, and A. Schröder, “Two-pulse 3D particle tracking with Shake-The-Box,” *Exp. Fluids* **64**, 93 (2023).
- ³⁴B. Wieneke, “Iterative reconstruction of volumetric particle distribution,” *Meas. Sci. Technol.* **24**, 024008 (2013).
- ³⁵J. Westerweel and F. Scarano, “Universal outlier detection for PIV data,” *Exp. Fluids* **39**, 1096–1100 (2005).
- ³⁶D. Schanz, S. Gesemann, and A. Schröder, “Shake-The-Box: Lagrangian particle tracking at high particle image densities,” *Exp. Fluids* **57**, 70 (2016).
- ³⁷A. Sciacchitano, B. Leclaire, and A. Schröder, “Main results of the first Lagrangian particle tracking challenge,” in 14th International Symposium on Particle Image Velocimetry (ISPIV 2021), 2021.
- ³⁸I. P. Christiansen, “Numerical simulation of hydrodynamics by the method of point vortices,” *J. Comput. Phys.* **13**, 363–379 (1973).
- ³⁹G. Bebis and M. Georgiopoulos, “Feed-forward neural networks,” *IEEE Potentials* **13**, 27–31 (1994).
- ⁴⁰I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, 2016) <http://www.deeplearningbook.org>.
- ⁴¹A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind, “Automatic differentiation in machine learning: A survey,” *J. Mach. Learn. Res.* **18**, 5595–5637 (2018).
- ⁴²D. P. Kingma and J. L. Ba, “Adam: A method for stochastic optimization,” in *the International Conference on Learning Representations (ICLR), San Diego, CA* (2014).
- ⁴³D. C. Liu and J. Nocedal, “On the limited memory BFGS method for large scale optimization,” *Math. Program.* **45**, 503–528 (1989).
- ⁴⁴S. Li and X. Feng, “Dynamic weight strategy of physics-informed neural networks for the 2D Navier–Stokes equations,” *Entropy* **24**, 1254 (2022).
- ⁴⁵Z. Xiang, W. Peng, X. Liu, and W. Yao, “Self-adaptive loss balanced physics-informed neural networks,” *Neurocomputing* **496**, 11–34 (2022).
- ⁴⁶S. Perez, S. Maddu, I. F. Sbalzarini, and P. Poncet, “Adaptive weighting of Bayesian physics informed neural networks for multitask and multiscale forward and inverse problems,” *J. Comput. Phys.* **491**, 112342 (2023).
- ⁴⁷P. G. Saffman, *Vortex Dynamics* (Cambridge University Press, 1992).
- ⁴⁸J. C. R. Hunt, A. A. Wray, and P. Moin, “Eddies, stream, and convergence zones in turbulent flows,” in Center for Turbulence Research Report CTR-S88, 1988.
- ⁴⁹J. F. G. Schneiders, I. Azizli, F. Scarano, and R. P. Dwight, “Pouring time into space,” in 11th International Symposium on Particle Image Velocimetry (PIV15), 2015.
- ⁵⁰R. J. Adrian, “Particle-imaging techniques for experimental fluid mechanics,” *Annu. Rev. Fluid Mech.* **23**, 261–304 (1991).