

Assimilating sparse 3D-PTV data of a pulsed jet using physics-informed neural networks

Ben Steinfurth^{1,*}, Abdelrahman Hassanein², N. Anh Khoa Doan², Fulvio Scarano²

1: Institute of Aeronautics and Astronautics, TU Berlin, Germany

2: Faculty of Aerospace Engineering, TU Delft, The Netherlands

* Correspondent author: ben.steinfurth@tu-berlin.de

Keywords: Particle tracking velocimetry, Data assimilation, Deep learning, Vortex rings

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 yields a higher spatial resolution, and the process remains robust, even at low particle densities ($ppp < 10^{-3}$). This is achieved by enforcing compliance with the governing equations, thus leveraging the spatio-temporal evolution of the measured flow field. The results indicate that the PINN reconstructs unambiguously velocity, vorticity and pressure fields 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 support the findings of Clark di Leoni (2023) suggesting that the PINN methodology is inherently suited to the assimilation of PTV data, in particular under conditions of severe sparsity or during experiments with limited control of seeding concentration.

1. 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 (e.g., Agüi & Jiménez, 1987). The most straightforward method is to collect velocity vectors into bins that yield the local ensemble-averages as demonstrated by Kasagi & Nishino (1991) and Agüera et al. (2016). Alternatively, linear interpolation circumvents the need to select a bin size, while adaptive Gaussian windowing (Agüi & Jiménez, 1987) and Gaussian radial basis functions have been shown to yield a continuously smooth reconstruction even in the presence of measurement noise (Vedula & Adrian, 2005; Casa & Krueger, 2014). 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 modelling of the underlying physical process.

In the past decade, methods have emerged that impose specific constraints based on the fluid dynamic governing equations. For instance, mass conservation can be imposed by vanishing velocity divergence during CGR (de Silva et al., 2013; Schiavazzi et al., 2014). In addition, the momentum equation can be used for the CGR 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+, Schneiders et al., 2014; Schneiders & Scarano, 2016; VIC#, Jeon et al., 2022; VIC-TSA, Scarano et al., 2022). Similarly, penalization of divergence as well as of a residual pertaining to the momentum equation have led to the FlowFit method, introduced by Gesemann et al. (2016) and further developed by Ehlers et al. (2020). Although these methods succeed in the dense reconstruction of velocity and vorticity fields from experiments performed at various levels of the seeding density, the integration of boundary conditions both at the edges of the domain (open boundaries) and at fluid-solid boundaries is regarded as problematic.

Alternative to the above methods of PTV data assimilation is the use of artificial neural networks, first attempted by Labonté (2001) to track particles across two image frames, obtaining the underlying velocity field¹⁴. 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 by Du et al. (2023). The PINN methodology has been introduced by Raissi et al. (2019, 2020) who identified several relevant aspects, in particular the incorporation of prior system knowledge by enforcing consistency with the underlying physics. Examples of this method to fluid mechanics are provided in a recent review by Cai et al. (2022), including successful solutions of inverse problems where unknown, or hidden, flow properties are assimilated.

Recently, Clark di Leoni et al. (2023) employed PINNs to infer velocity and pressure fields from numerical and experimental PTV data, comparing the results with the Constrained Cost Minimization method introduced by Agarwal et al. (2021). The objective in the present article is to assess the capability of PINNs in assimilating experimental 3D-PTV data. Most of the above works require the use of numerically simulated experiments in order to apply error metrics to a ground truth solution. In the present study, 3D PTV measurements of a pulsed jet at $Re \approx 3000$ are conducted with a variable density 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. In the present case, the experimental ground truth is produced with

additional high-resolution planar PIV measurements. The experimental apparatus and measurement procedure are introduced first, followed by an overview of assimilation methods compared in this paper (section III). Then, in section 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 section V. The main learnings and conclusions are provided in section VI.

2. Layout of experiments

2.1. Pulsed jet facility and operating 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 (Steinfurth & Weiss, 2020, 2021, 2022).




Figure 1 Pulsed-jet actuator used to produce axisymmetric vortex rings

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 millimeters upstream of the outlet. Under some assumptions (Steinfurth, 2022), 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 $u_{\text{jet}} = 10$ m/s, translating into a Reynolds number of $Re = u_{\text{jet}}D/\nu = 3300$. Considering the concept of

a formation number (Gharib et al., 1998), maximum-circulation vortex rings can be expected at $L/D \approx 4$. In the present study, L is interpreted as the time-dependent length of a virtual fluid column emerging from the nozzle exit at bulk velocity, $L = tu_{\text{jet}}$. This suggests that the vortex ring formation is completed at $t = 2$ 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.

The dynamical evolution of the jet and its vortex ring is illustrated in Fig. 2 as obtained with high-resolution planar 2C-PIV measurements. The appearance and growth of the axisymmetric vortex is highlighted with the out-of-plane (viz. azimuthal) vorticity. Note that the nondimensional 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 maximum axial velocity on the center line and the time coordinate, where $t = 0$ s marks the moment when the jet emerges from the exit plane ($x = 0$ mm). At $t^* = 1.6$, a first vortex core detaches from the wall, followed by a second one at $t^* = 4$. 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.

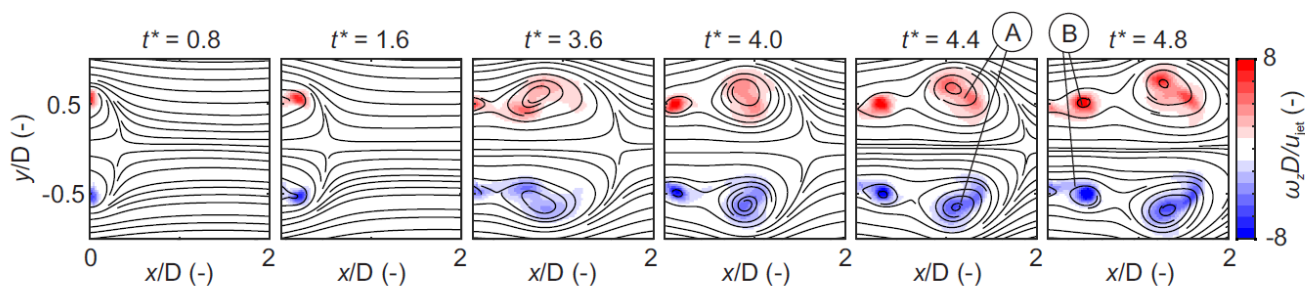


Figure 2 Vortex ring formation: out-of-plane vorticity overlaid with streamlines in vortex ring frame of reference

The measurements are performed seeding the air jet with micron-sized DEHS tracer particles (Kähler et al., 2002), supplied to the compressed air line feeding the jet. The tracer concentration was controlled by varying the flow through the Laskin nozzle-type particle generator. A secondary, unseeded, air feed is connected to the supply for a controlled dilution of the seeded air. The jet exhausts inside a confined chamber with acrylic glass walls of dimensions 500 mm (axial) \times 250 mm (radial) such as to homogenize the concentration of the seeding particles from the jet to the quiescent ambience. The particles were illuminated across the jet symmetry plane with an EverGreen200 dual-cavity Nd:YAG laser ($\lambda = 532\text{nm}$, 2×200 mJ pulse energy). A light sheet with a maximum thickness of $0.1D$ (500 μm) 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, $6.5 \mu\text{m}$ pixel pitch) equipped with an $f = 105$ mm objective and an fstop 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 acquisition to the signal controlling the magnetic valve. The time separation of 12ms corresponds to a particle displacement of 0:12mm (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 96×96 px² down to 16×16 px², the latter of which corresponds to a spatial resolution of $\Delta x = \Delta y = 0.2$ mm. An overlap factor of 75% reduces the vector pitch to approximately 0:05 mm ($0.01D$).

2.2. 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.




Figure 3 PTV setup; detailed view in the top left shows how the original light sheet highlighted by the green

The seeding procedure is the same as that for the planar PIV measurements, however, the tracer concentration was set to a significantly lower value to mitigate the occurrence of ghost particles. The laser beam of 7mm diameter was expanded into an elliptical cross section of approximately 40mm by 20mm 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 (Ghaemi

& Scarano, 2010), 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, the same as was used for the planar experiment, is placed subtending a tomographic aperture of 50 degrees. The f-stop 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 procedure (Wieneke, 2008) reduces residual calibration errors below 0.1 px.

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 ms ($\Delta t^* = 0.2$) in the range $t^* = [0; 4]$ and 200 ms ($\Delta t^* = 0.4$) in the range $t^* = [0; 4]$. While approximately constant for each phase (measurement set), the seeding density varied in the range $ppp = 10^{-4} - 10^{-3}$ 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.




Figure 4 Representative particle images for three phases; top: raw images, bottom: pre-processed images

The 3D particle detection and motion analysis was performed using the two-pulse variant (Novara et al., 2023) of the shake-the-box (STB) 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 1 px and applying three iterations of particle reconstruction and tracking (Wieneke, 2013) To remove data outliers, a spatial median filter (Westerweel & Scarano, 2005) was applied after particle tracking. Example particle distributions obtained at three different seeding conditions (the same as in Fig. 4) are shown in Fig. 5. Outliers, characterized by a streamwise velocity exceeding u_{jet} , are included in the top row whereas the images in the bottom row show the filtered PTV results that are used for the data assimilation.

The three-dimensional analysis returns the particle positions in physical space. The STB algorithm performs the pairing and determines the velocity of each tracer. Fig. 5 illustrates the instantaneous distribution of particle pairs, color-coded by axial velocity.

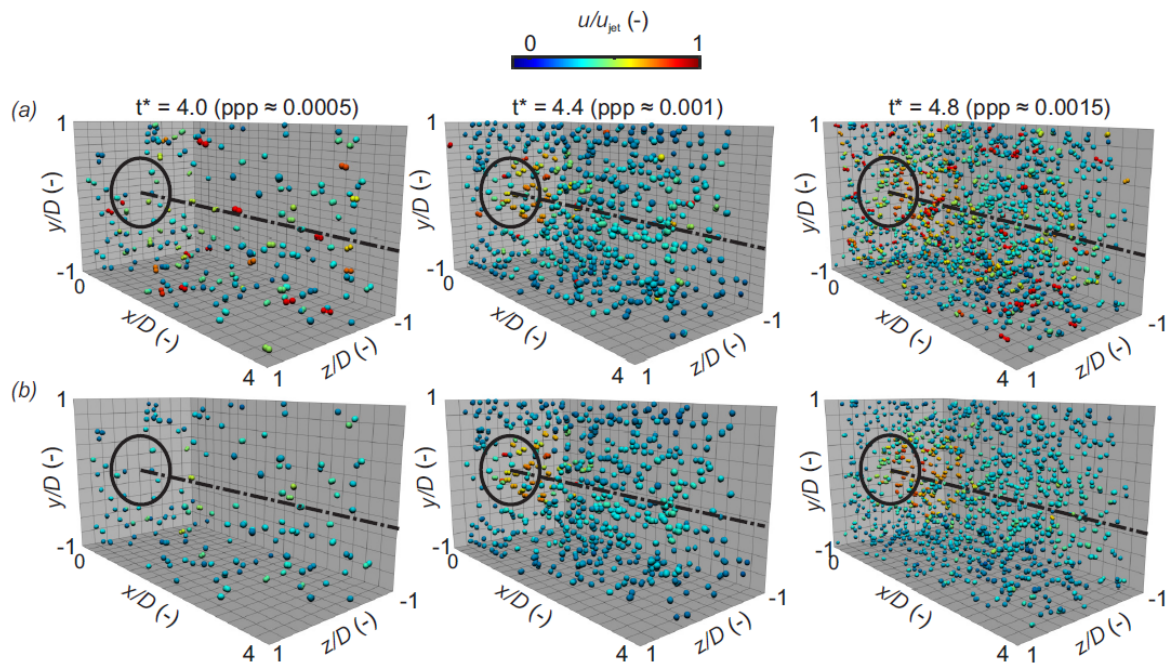


Figure 5 Examples of particle pairs obtained at different phases and seeding levels (left to right). Top row: includes outliers. Bottom row: after data validation; jet outlet and axis indicated by ellipse and dash-dotted line, respectively

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 $ppp \approx 10^{-3}$ (right-hand side in Fig. 5) is still far below the concentration of $ppp = 0.1$ that can be handled by the STB algorithm (Schanz et al., 2016; Sciacchitano et al., 2021). 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 5 pixels, which lowers the upper limit for the concentration.

In order to increase the effective spatial concentration of the velocity field, tracks were accumulated from 100 snapshots taken at the same phase of the pulsed jet. The approach was verified considering that the cycle-to-cycle velocity fluctuations did not exceed 5% of the jet exit velocity (see Appendix A). The number of detected particle tracks at each phase is shown in Fig. 6, yielding significant variations depending on the phase, from a minimum of 10,000 up to 60,000. Three cases are selected for further analysis that span such range and are separated by a relatively small time step (highlighted by red circles), which correspond to the time steps presented in Figs. 4 and 5.




Figure 6 Average number of particle tracks and particle image density per snapshot as a function of pulsed jet phase; error bars span one standard deviation; red circles highlight the conditions considered in the present article

3. 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 $(10 \times 10 \times 10)$ mm³.

3.1. 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 (Vedula & Adrian, 2005). 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, although the scaling is rather unfavourable (to halve the bin linear size one needs an eight-fold increase of the particle concentration).

In the present study, cubic bins with edge lengths of 1 mm and an overlap of 75% were chosen, and a second-order polynomial was used as fitting function inside the bins. As a result, the data are represented on a grid of $40 \times 40 \times 40$ points spaced by 0.25 mm.

3.2. Vortex-in-cell assimilation

More advanced CGR methods follow the vortex-in-cell (VIC) paradigm (Christiansen, 1973) by considering the governing equations in vorticity-velocity formulation. While initially used to enhance the temporal resolution of tomographic PIV data (Schneiders et al., 2014), the method was proven to be applicable to leverage time-resolved recordings for spatial interpolation of scattered data (Schneiders & Scarano, 2016). 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 turbulent dissipation rate with reasonable accuracy (Schneiders & Scarano, 2016). Its main drawback lies in a sensitivity towards data at the boundaries of the domain, which was recently addressed introducing some ingenious modifications (VIC#, Jeon et al., 2022).

In the present study, the VIC# technique, implemented in the *LaVision DaVis 10* software, is employed to reconstruct the velocity field on a $40 \times 40 \times 40$ grid (resolution: ca. 0.25 mm) by applying 40 iterations per snapshot. It is important to mention that only data from individual phases are taken into account in the present study. Hence, only velocities (but no material acceleration) are provided to the algorithm.

3.3. 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. 7). Note that a tilde indicates normalization to the range

[0;1] to avoid training issues that could arise from network input/output differing by orders of magnitude.




Figure 7 Schematic representation of PINN; orange: input coordinates (that can be chosen arbitrarily), purple: hidden layers, red: output quantities, green: automatic differentiation

A multilayer perceptron, also called feedforward neural network (Bebis & Georgiopoulos, 1994), 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 $\boldsymbol{\zeta}_l = \mathbf{W}_l^T \boldsymbol{\xi}_{l-1} + \mathbf{b}_l$ where $\boldsymbol{\xi}_{l-1}$ is the output of the previous layer and \mathbf{b}_l is the bias in layer l (Goodfellow et al., 2016). Finally, non-linear behavior is introduced through the activation g , as $\boldsymbol{\xi}_l = g(\boldsymbol{\zeta}_l)$. 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 the next section. No further weight is required since the minimization of LBC 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:

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

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 $\mathbf{u}_{\text{PTV}}(\mathbf{x}_i, t_i)$ are available. The prediction error \mathcal{L}_{PTV} therefore allows the network to anchor its output at the locations where information is available.

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

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

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

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

In the above equation, the required gradients of \mathbf{u} and p with respect to space and time are readily obtained using automatic differentiation (Baydin et al., 2018). The air density and kinematic viscosity are set to $\rho = 1.25 \text{ kg/m}^3$ and $\nu = 1.5e^{-5} \text{ m}^2/\text{s}$, respectively, corresponding to laboratory conditions during experiments. It is worth mentioning that the PINN function (Eq. 1) can be evaluated at spatio-temporal locations other than those of the training data and therefore, the collocation points can be arbitrarily set 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 output grid in the VIC method (Schneiders & Scarano, 2016). Furthermore, the

space-time coverage of the PINN method can be regarded as an equivalent of the time segment optimization in the VIC-TSA technique (Scarano et al., 2022).

The third term in Eq. 2, \mathcal{L}_{BC} , ensures that the PINN output satisfies boundary conditions that are represented by vanishing velocity and a base value of zero pressure at a transverse jet center line 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 $u = 0$ m/s and $p = 0$ Pa when scaled to the range [0;1]):

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

The overall loss (Eq. 2) was minimized using the ADAM optimizer³⁸ at a learning rate of $lr = 1e^{-3}$ for 1000 epochs. Subsequently, to further finetune the network, 15,000 evaluations of the loss function were handled by a Limited-memory Broyden-Fletcher-Goldfarb-Shanno (LBFGS) optimizer (Liu & Nocedal, 1983) driving the supremum norm of the loss function gradient below $\|\nabla\mathcal{L}\|_{\infty} = 1e^{-8}$ for all cases presented in the next section.

4. 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 the following subsections, the design choices with regards to these parameters are discussed.

4.1. 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 towards the number of collocation points N_c per timestep is illustrated in Fig. 8. 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 phase N_t will be provided later. Clearly, the output is altered substantially for $N_c/N_t < 1000$ whereas not much is to be gained for larger numbers of collocation points. Since N_c drives the computational cost associated with the PINN training, a moderate number of $N_c/N_t = 2000$ is chosen.




Figure 8 Influence of number of collocation points per training phase on velocity predictions

4.2. 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. 9, 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 $\mathcal{R} = \nabla \mathbf{u} + \partial_t \mathbf{u} + \mathbf{u} \nabla \mathbf{u} + (\nabla p)/\rho - \nu \nabla^2 \mathbf{u}$ 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 enforced with weights of $\lambda_{\text{NS}} = 1e^{-9}$ (center row) and $\lambda_{\text{NS}} = 1e^{-7}$ (bottom row), respectively.

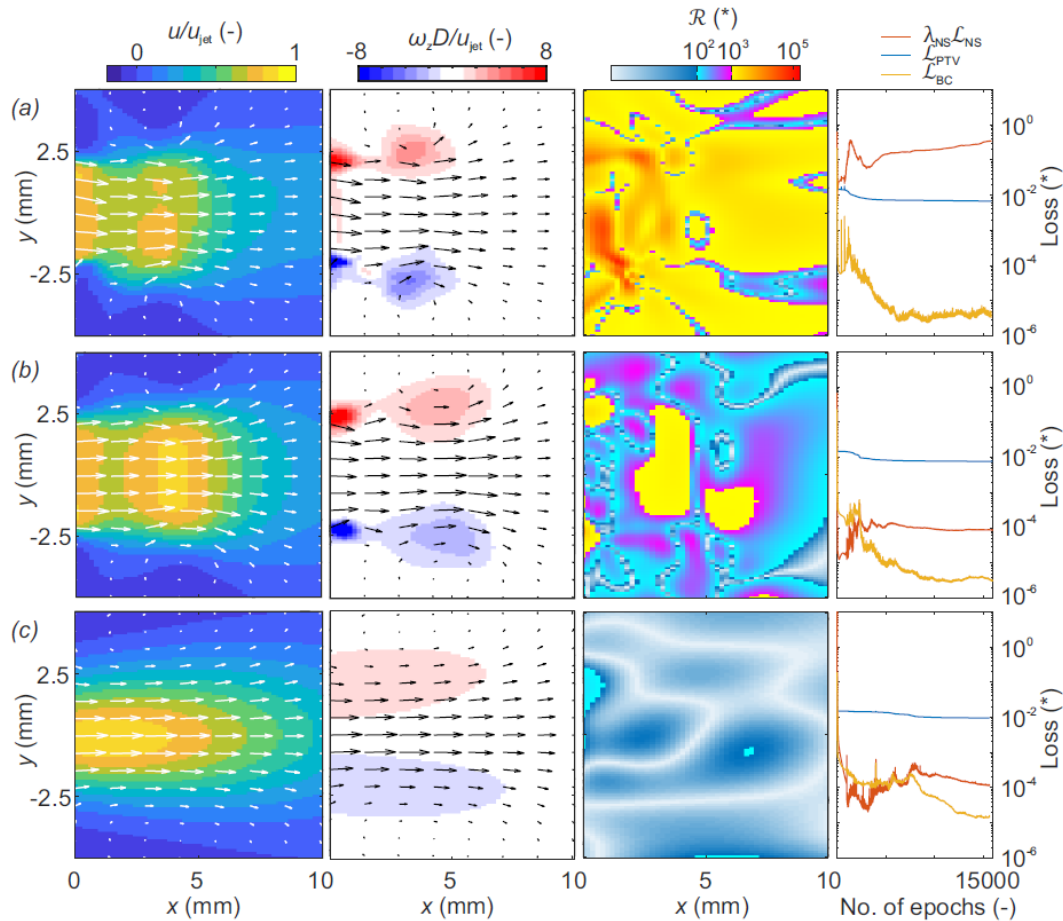


Figure 9 Influence of physics-informed loss on model output; (a) Navier-Stokes equations not taken into account during training, (b) physics-based loss weighted with $\lambda_{\text{NS}} = 1e^{-9}$ and (c) $\lambda_{\text{NS}} = 1e^{-7}$

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 yields a consistent velocity field, yet the vorticity pattern is largely 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-Stokes equations. The loss terms development during training (right column) shows a

monotonic decrease of LPTV while the physics-based term increases. In conclusion, the unconstrained optimization increasingly violates the Navier-Stokes equations.

In the center row of Fig. 9, a weight of $\lambda_{\text{NS}} = 1e^{-9}$ 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.

Further enforcing compliance to the Navier-Stokes equations ($\lambda_{\text{NS}} = 1e^{-7}$, bottom row), the physics-based loss becomes comparable to $\lambda_{\text{NS}} = 1e^{-9}$ (i.e., $\mathcal{O}(\mathcal{L}_{\text{NS}}) = 10^4$), which implies that the residuals are 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 towards the trivial (i.e. homogeneous) solution of the Navier-Stokes equations, which greatly departs from the measurement data.

By increasing the contribution of the physics-based loss, the influence of the PTV data is marginalized, allowing for larger deviations between the measurement data and the PINN output. This conflict is illustrated in Fig. 10, 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.

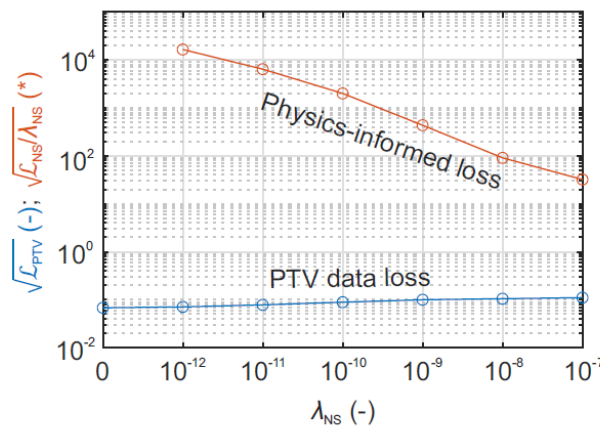


Figure 10 Loss terms as a function of the physics-based weight coefficient

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 (Li & Feng, 2022; Xiang et al., 2022; Perez et al., 2023), a heuristic approach is proposed based upon a crossed evaluation of the PINN output (e.g., Fig. 9) and the data loss term (Fig. 10). The model output may be required to follow the measurement data within a reasonable threshold conservatively estimated in the order of 10% of the reference velocity. This criterion would yield a weight $\lambda_{NS} = 1e^{-9}$, corresponding to the condition of optimal PINN output, based on the analysis of the flow field.

4.3. Number of training timesteps

When designing a PINN for dense flow field reconstruction, a choice needs to be made regarding the length of the modelled time series. For instance, one may use a single PINN for the entire available time series. However, this would necessitate a very expressive network that requires extensive training time. 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 timestep 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 (15000) is shown in Fig. 11 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.




Figure 11 Influence of number of training phases on velocity, vorticity and pressure predictions after the same number of epochs; (a) $N_t = 3$, (b) $N_t = 5$, (c) $N_t = 7$

Clearly, the training for cases b and c has not converged, which is manifested in blurred spatial velocity gradients yielding an elongated shear-layer where the primary and secondary vortex ring become 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 timesteps as the collocation points are sampled randomly inside the range spanned by the preceding and succeeding training

phases. The training time for $N_t = 3$ phases was approximately 15 minutes (using a V100 Tensor Core GPU).

4.4. Network depth

Finally, the width and depth of the PINN need to be chosen. Keeping the number of neurons per layer constant at 50 (each activated with a \tanh function), the number of layers was varied and the overall loss evaluated (Fig. 12).




Figure 12 Overall loss as a function of the number of hidden layers; training with $N_t = 3$ phases and target phase $t^* = 4.0$

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.

5. Demonstration of PTV data assimilation

In this section, the PINN method is compared to alternative CGR techniques, namely data binning and VIC#. It is important to mention that only for the PINN method, measurement data from different phases is taken into account as introduced above.

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 $x = [0,10]$ mm is on the order of $\mathcal{O}(N_p) = 10^4$, which was facilitated by accumulating particle tracks 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 $\mathcal{O}(ppp) = 10^{-1}$ (accumulated particles, highest concentration) and $\mathcal{O}(ppp) = 10^{-5}$ (single snapshot, lowest concentration).

5.1. 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. 13. Recall that for the binning method, cubes with an edge length of 1 mm (75% overlap) were chosen, yielding a vector pitch of 0.25 mm. The same spatial resolution is imposed with the VIC# method. The PINN output can be inferred at arbitrary coordinates and it was sampled with 0.1 mm grid spacing.

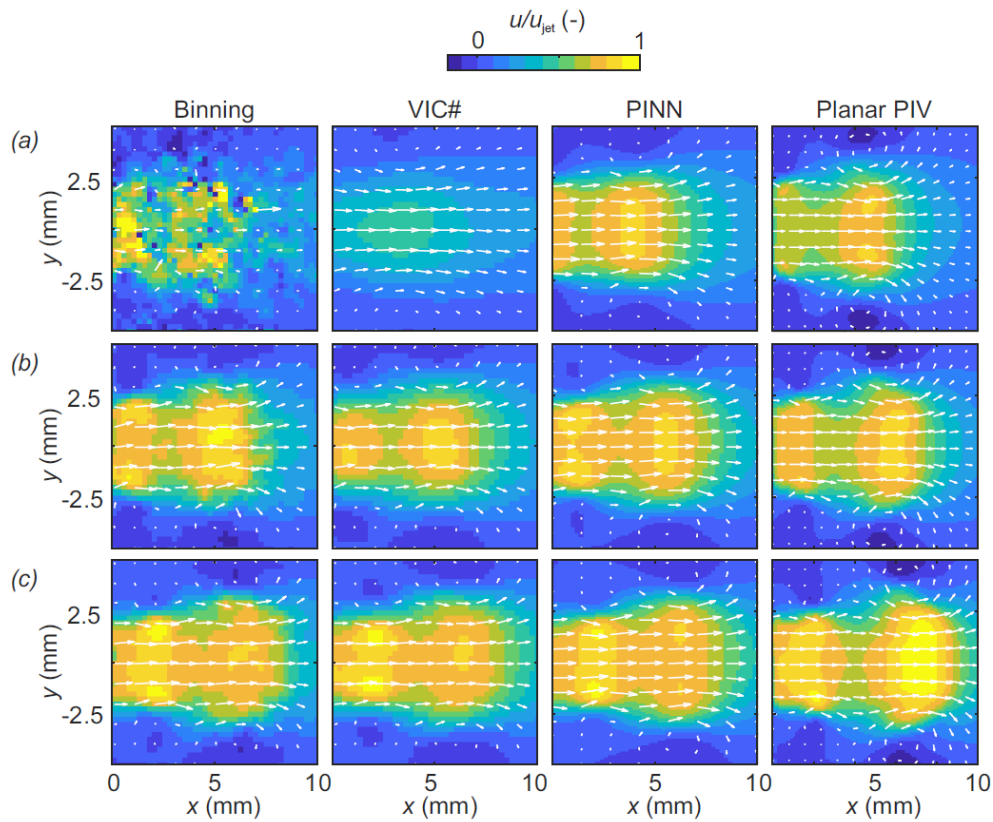


Figure 13 Velocity fields obtained by different methods compared to reference planar PIV results, (a) $t^* = 4.0$ ($N_p = 6700$), (b) $t^* = 4.4$ ($N_p = 27300$), (c) $t^* = 4.8$ ($N_p = 59000$)

For the first phase ($t^* = 4.0$, $ppp = 2 \cdot 10^{-2}$), the binning of approximately 6700 velocity vectors leads to a discontinuous velocity field where the region spanned by the jet is barely recognized. The VIC# method yields a smoother distribution but the axial velocity is considerably underpredicted. The PINN method 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 timesteps, featuring a higher seeding concentration, a clearer picture is revealed by the binning method. Nonetheless, there remain distinct discontinuities that are not present in the VIC# and PINN output that feature a higher degree of spatial coherence.

The above discussion is extended to the accuracy and resolution of the spatial velocity gradient by observing the reconstructed out-of-plane vorticity (Fig. 14).

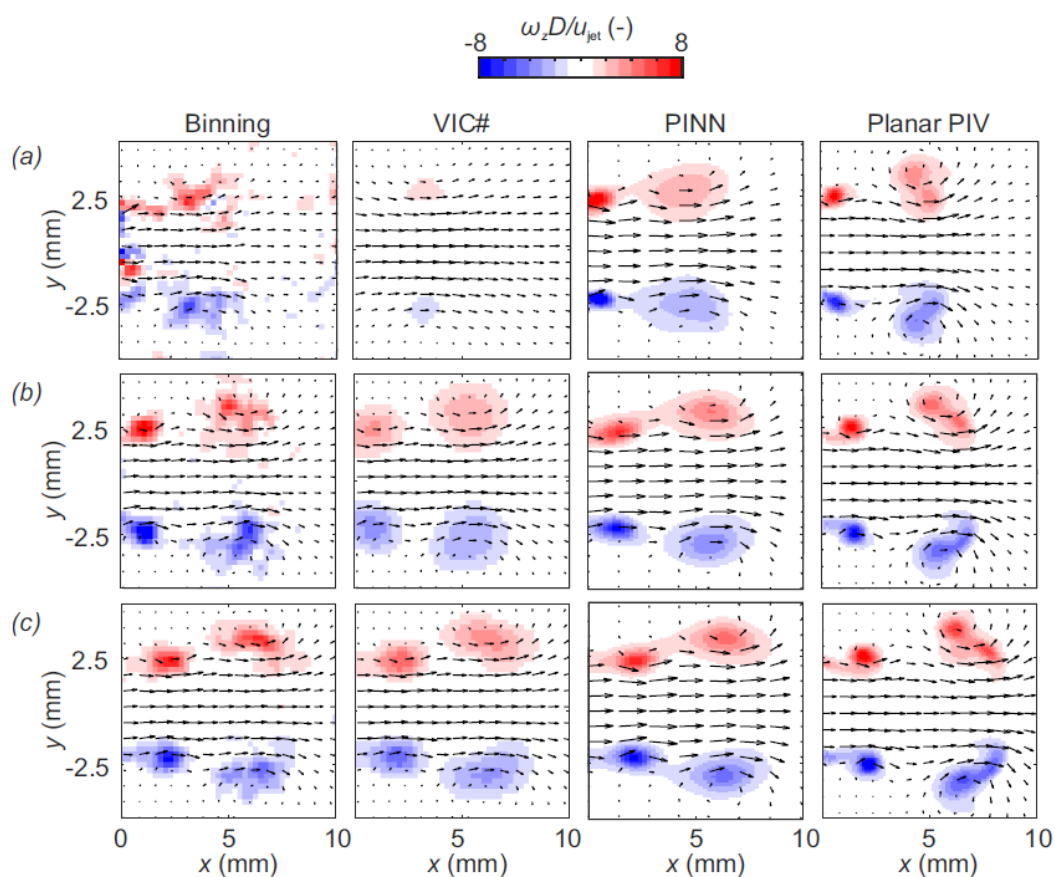


Figure 14 Vorticity fields obtained by different methods compared to reference planar PIV results, (a) (b) and (c) same as in Fig. 13

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 unrewarding at $t^* = 4.0$, the binning method only produces noisy vorticity distributions at the subsequent phases where larger numbers of particle tracks are available. As for the VIC# and PINN predictions, the latter yields a

stronger degree of similarity with the reference PIV measurements. Yet, it is unable to capture the inner structure of the detached vortex ring, which seems to be composed of two merged vortices.

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 (Saffman, 1992):

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

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, (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

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

is obtained for different timesteps. Note that this quantity decreases as the vortex ring grows in thickness during its formation (Fig. 15).




Figure 15 Development of non-dimensional vortex ring energy obtained by different methods compared to reference planar PIV results

The limiting value reported by Garib et al. (1998) 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 $t^* > 5.0$. Whereas a deviation is noticed at smaller formation times, the α curve obtained for the PINN output almost collapses with the planar PIV data otherwise. In contrast, more substantial

differences are on display for the binning and VIC# methods where the non-dimensional energy does not drop below $\alpha = 0.4$.

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. 16, the full three-dimensional structure of the flow field is represented by means of iso-surfaces.




Figure 16 Time series of iso-contours used for vortex detection obtained by different methods, (a) $t^* = 4.0$ ($N_p = 6700$), (b) $t^* = 4.4$ ($N_p = 27300$), (c) $t^* = 4.8$ ($N_p = 59000$): Q criterion shown at $Q = 5e^6 s^{-1}$ (red) and low-pressure region at $\Delta p = -11$ Pa (blue)

The Q criterion introduced by Hunt et al. (1988) (red iso-surface) is adopted to unambiguously detect the vortex ring produced by the pulsed jet. Furthermore, the PINN solution 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 to the vortex core and it has been often considered as equivalent vortex identification criterion⁴⁴. 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 $t^* = 4.0$ where the VIC# method fails due to severe data sparsity. Furthermore, the formation of the secondary vortex ring can be observed which is caused by the impulsive ejection of fluid associated with the trailing

jet. As for the binning results, noisy vortex structures are only perceptible at $t^* = (4.4, 4.8)$ whereas no distinct vortex ring can be observed for the first phase.

In the pressure field predicted by the PINN, toroidal structures corresponding to the primary and secondary vortex rings are adequately captured at $t^* = (4.4, 4.8)$ as 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.

5.2. Training on single snapshots

In the following, even more challenging conditions for data assimilation are considered. Although 3D PTV is frequently performed with time-resolved measurements, experiments at higher flow velocity need to revert to the double-frame mode, such as in the present case. It is therefore relevant to compare the behaviour of CGR methods applied to single snapshots of relatively sparse data (Fig. 17).

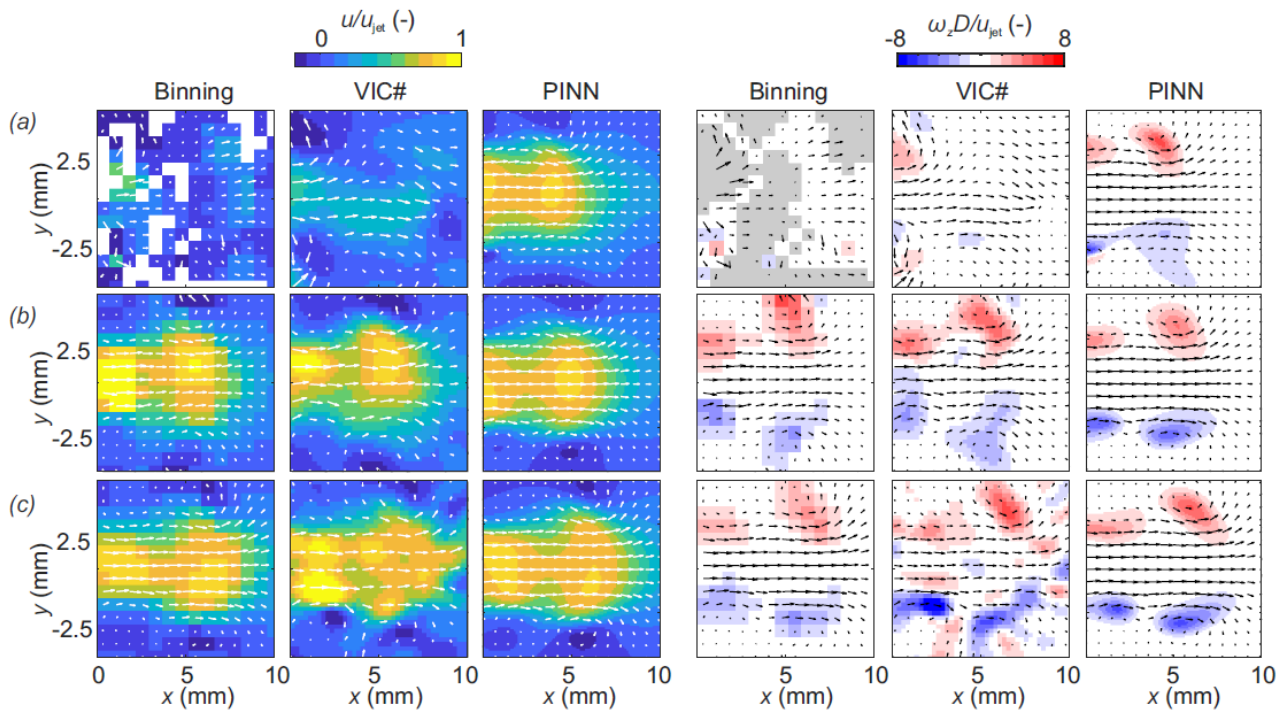


Figure 17 Velocity and vorticity fields obtained by different methods; (a) $t^* = 4.0$ ($N_p \approx 100$), (b) $t^* = 4.4$ ($N_p \approx 400$), (c) $t^* = 4.8$ ($N_p \approx 700$)

At the lowest particle concentration ($N_p = 89$, i.e. less than 0.1 particles per mm³), a substantial performance gap is observed at the first phase. For the binning method, no data are available

inside an appreciable proportion of the symmetry plane, and the flow is barely recognizable otherwise. While filling these gaps, the VIC# method again underpredicts the velocity magnitude as was the case with the accumulated particle tracks (Fig. 13). The PINN, on the other hand, predicts a velocity distribution that is similar to the reference velocity field shown in Fig. 13, clearly benefitting 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# method 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 timesteps. However, the largest similarity with the true velocity field shown in Fig. 13 is achieved by the PINN. This also applies to the vorticity fields where the binning method suffers from low spatial resolution and the VIC# yields incoherent results. To the contrary, 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. 14) is relatively small.

Finally, the same type of iso-contours as in Fig. 16 are presented for the evaluation of single snapshots Fig. 18.

As was the case for the velocity and vorticity fields presented above, there are significant performance drop-offs for the binning and VIC# methods when the data sparsity is increased. Both methods produce very 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 much clearer vortex identification. Even for the first snapshot, parts of this flow structure can be identified.

To summarize this section, we conclude that assimilating PTV data using PINNs is superior to the binning method and data assimilation using VIC#. However, it must be reiterated that for the latter, only one timestep is taken into account in the present study. Nonetheless, it appears that smaller amounts of velocity data are required to infer high-resolution velocity fields of larger spatial coherence with the PINN method.




Figure 18 The same as in Fig. 16 but now, iso-surfaces obtained from a single snapshot; (a) $t^* = 4.0$ ($N_p \approx 100$), (b) $t^* = 4.4$ ($N_p \approx 400$), (c) $t^* = 4.8$ ($N_p \approx 700$)

6. 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 Clark di Leoni et al. (2023). 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 defined by Cartesian coordinates and time. This mapping was learned by constraining the PINN output by short time series of measurement data on the one hand and the three-dimensional, incompressible Navier-Stokes equations on the other. One training cycle comprising 16,000 evaluations of the loss function took approximately 15 minutes (on a V100 Tensor Core GPU).

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. Disregarding the PTV data, in contrast, yields physics-compliant, but 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 measurement uncertainty can be applied as a reasonable threshold for the latter.

Following this approach, we found that the constructed PINNs are capable of predicting reliable velocity fields of high spatial coherence. Moreover, they are well-suited to handle the reduction of scattered PTV data onto Cartesian grids seamlessly since arbitrarily structured data can be provided for training and inference. In this regard, a comparison was carried out with two alternative standard methods, namely binning and VIC#. Since the particle concentration was relatively low, bin dimensions that were large compared to characteristic lengths of the flow needed to be chosen, yielding a low spatial resolution. Although VIC# performs reasonably well in cases of higher particle numbers, it fails to overcome the data sparsity otherwise. The same is not true for the PINN that is shown to adequately replicate the pattern underlying the particle-based measurements. Consequently, a clearer picture of the occurring flow 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 despite a seeding density which does not exceed $ppp = 10^{-3}$. In addition, the PINN allows to infer regions of low pressure, assimilated through the Navier-Stokes equations, that can augment the vortex identification.

The superiority of PINNs compared to the alternative methods assessed in this study can be attributed to two properties. First, the physics-based approach enables de-noising of the measurement data but, perhaps more importantly, it inherently promotes a differentiable output resulting in more reliable quantities of interest that are derived from the velocity field. Second, data sparsity is managed more robustly by the physics-based spatio-temporal interpolation, which is consistent with the concept of 'pouring time into space' (Schneiders et al., 2015). This feature may be leveraged in order to reduce the amount of data required to obtain sufficient information of the flow. Looking at this from a different perspective, PINNs can be applied to assimilate PTV data acquired under non-ideal experimental conditions leading to a low particle concentration.

In summary, considerable proof has been presented that PINNs represent a tool well-suited for the assimilation of PTV data. Future studies may be directed at testing this technique under different boundary conditions. Furthermore, effort should be dedicated to establishing a non-heuristic method of balancing the loss function.

Acknowledgements

The authors gratefully acknowledge financial support from the Deutscher Akademischer Austauschdienst (DAAD, German Academic Exchange Service) - under program ID 57664191.

References

- Agarwal, K., Ram, O., Wang, J., Lu, Y. and Katz, J. (2021) Reconstructing velocity and pressure from noisy sparse particle tracks using constrained cost minimization. *Experiments in Fluids*
- Agüera, N., Cafiero, G., Astarita, T., and Discetti, S. (2016) Ensemble 3D PTV for high resolution turbulent statistics. *Measurement Science and Technology*. 27
- Agüi, J. C. and Jiménez, J. (1987). On the performance of particle tracking. *Journal of Fluid Mechanics*, 185, 447– 468
- Baydin, A. G., Pearlmutter, B. A., Radul, A. A. and Siskind, J. M. (2018) Automatic differentiation in Machine Learning: A survey. *Journal of Machine Learning Research*. 18. 5595–5637
- Bebis, G. and Georgiopoulos, M. (1994) Feed-forward neural networks. *IEEE Potentials*. 13. 27–31
- Goodfellow, I., Bengio, Y. and Courville, A. (2016) *Deep Learning* (MIT Press, 2016)
- Cai, S., Mao, Z., Wang, Z., Yin, M. and Karniadakis, G. E. (2022) Physics-informed neural networks (PINNs) for fluid mechanics: A review,” *Acta Mechanica Sinica*. 37. 1727–1738
- Casa, L. D. C. and Krueger, P. S. (2014) Radial basis function interpolation of unstructured, three-dimensional, volumetric particle tracking velocimetry data. *Measurement Science and Technology*. 24
- Christiansen, I. P. (1973) Numerical simulation of hydrodynamics by the method of point vortices. *Journal of Computational Physics*. 13. 363–379
- Clark Di Leoni, P., Agarwal, K., Zaki, T. A., Meneveau, C. and Katz, J. (2023) Reconstructing turbulent velocity and pressure fields from under-resolved noisy particle tracks using physics-informed neural networks. *Experiments in Fluids*. 64
- de Silva, C. M., Philip, J. and Marusic, I. (2013) Minimization of divergence error in volumetric velocity measurements and implications for turbulence statistics. *Experiments in Fluids*. 54
- Du, Y., Wang, M., and Zaki, T. A. (2023) State estimation in minimal turbulent channel flow: A comparative study of 4DVar and PINN. *International Journal of Heat and Fluid Flow*.
- Ehlers, F., Schröder, A. and Gesemann, S. (2020) Enforcing temporal consistency in physically constrained flow field reconstruction with FlowFit by use of virtual tracer particles. *Measurement Science and Technology*. 31

- Gesemann, S., Huhn, F., Schanz, D. and Schröder, A. (2016) From noisy particle tracks to velocity, acceleration and pressure fields using B-splines and penalties. *18th International Symposium on Applications of Laser Techniques to Fluid Mechanics*
- Ghaemi, S. and Scarano, F. (2010) Multi-pass light amplification for tomographic particle image velocimetry applications. *Measurement Science and Technology*. 21
- Gharib, M., Rambod, E. and Shariff, K. (1998) A universal time scale for vortex ring formation. *Journal of Fluid Mechanics*. 360. 121–140
- Hunt, J. C. R., Wray, A. A. and Moin, P. (1988) Eddies, stream, and convergence zones in turbulent flows in *Center for Turbulence Research Report CTR-S88*
- Jeon, Y. J., Müller, M. and Michaelis, D. (2022) Fine scale reconstruction (VIC#) by implementing additional constraints and coarse-grid approximation into VIC+. *Experiments in Fluids*. 63
- Kähler, C. J., Sammler, B. and Kompenhans, J. (2002) Generation and control of tracer particles for optical flow investigation in air. *Experiments in Fluids*. 33. 736–742
- Kasagi, N. and Nishino, K. (1991) Probing turbulence with three-dimensional particle-tracking velocimetry. *Experimental Thermal and Fluid Science*. 4, 601–612
- Kingma, D. P. and Ba, J. L. (1989) Adam: A method for stochastic optimization," *Open Journal of Statistics*. 11
- Li, S. and Feng, X. (2022) Dynamic weight strategy of physics-informed neural networks for the 2D Navier–Stokes equations. *Entropy*. 24
- Liu, D. C. and Nocedal, J. (1989) On the limited memory BFGS method for large scale optimization," *Mathematical Programming*. 45. 503–528
- Labonté, G. (2001) Neural network reconstruction of fluid flows from tracer-particle displacements. *Experiments in Fluids*. 30. 399-409
- Novara, M., Schanz, D. and Schröder, A. (2023) Two-Pulse 3D particle tracking with Shake-The-Box. *Experiments in Fluids*. 64
- Perez, S., Maddu, S., Sbalzarini, I. F. and Poncet, P. (2023) Adaptive weighting of Bayesian physics informed neural networks for multitask and multiscale forward and inverse problems. *Journal of Computational Physics*. 491
- Raissi, M., Perdikaris, P. and Karniadakis, G. E. (2019) Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*. 378. 686–707
- Raissi, M., Yazdani, A. and Karniadakis, G. E. (2020) Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Science*. 367. 1026–1030
- Saffman, P. G. (1992) *Vortex dynamics* (Cambridge University Press)

- Scarano, F., Schneiders, J. F. G., Saiz, G. G. and Sciacchitano, A. (2022) Dense velocity reconstruction with VIC-based time-segment assimilation. *Experiments in Fluids*. 63
- Schanz, D., Gesemann, S. and Schröder, A. (2016) Shake-The-Box: Lagrangian particle tracking at high particle image densities *Experiments in Fluids*. 57
- Schiavazzi, D., Coletti, F., Iaccarino, G. and Eaton, J. K. (2014) A matching pursuit approach to solenoidal filtering of three-dimensional velocity measurements. *Journal of Computational Physics*. 263. 206– 221
- Schneiders, J. F. G., Dwight, R. P. and F. Scarano (2014) Time-supersampling of 3D-PIV measurements with vortex-in-cell simulation. *Experiments in Fluids*. 55
- Schneiders, J. F. G., Azijli, I., Scarano, F. and Dwight, R. P. (2015) Pouring time into space. *11th International Symposium on Particle Image Velocimetry - PIV15*
- Schneiders, J. F. G. and Scarano, F. (2016) Dense velocity reconstruction from tomographic PTV with material derivatives. *Experiments in Fluids*. 57
- Sciacchitano, A., Leclaire, B. and Schröder, A. (2021) Main results of the first Lagrangian particle tracking challenge. *14th International Symposium on Particle Image Velocimetry – ISPIV2021*
- Steinfurth, B. and Weiss, J. (2020) Vortex rings produced by non-parallel planar starting jets. *Journal of Fluid Mechanics*. 903
- Steinfurth, B. and Weiss, J. (2021) Velocity ratio effect on flow structures of non-parallel planar starting jets in cross-flow. *Journal of Fluid Mechanics*. 915
- Steinfurth, B. and Weiss, J. (2022) Modelling the decay of finite-span starting and stopping wall jets in an external stream. *Journal of Fluid Mechanics*. 951
- Steinfurth, B. (2023) Flow physics of pulsed-jet actuation. Ph.D. thesis, TU Berlin
- Vedula, P. and Adrian, R. J. (2005) Optimal solenoidal interpolation of turbulent vector fields: application to PTV and super-resolution PIV. *Experiments in Fluids*. 213–221
- Westerweel, J. and Scarano, F. (2005) Universal outlier detection for piv data. *Experiments in Fluids*. 39. 1096– 1100
- Wieneke, B. (2008) Volume self-calibration for 3D particle image velocimetry. *Experiments in Fluids*. 45. 549– 556
- Wieneke, B. (2013) Iterative reconstruction of volumetric particle distribution. *Measurement Science and Technology*. 24
- Xiang, Z., Peng, W., Liu, X. and Yao, W. (2022) Self-adaptive loss balanced physics-informed neural networks,” *Neurocomputing*. 496. 11–34.

Appendix A: Reproducibility of jet flow

As a measure for the reproducibility of the flow, the mean fluctuating part of the main velocity component is presented in Fig. 19. Recall that for each phase, 100 instantaneous velocity fields are considered.

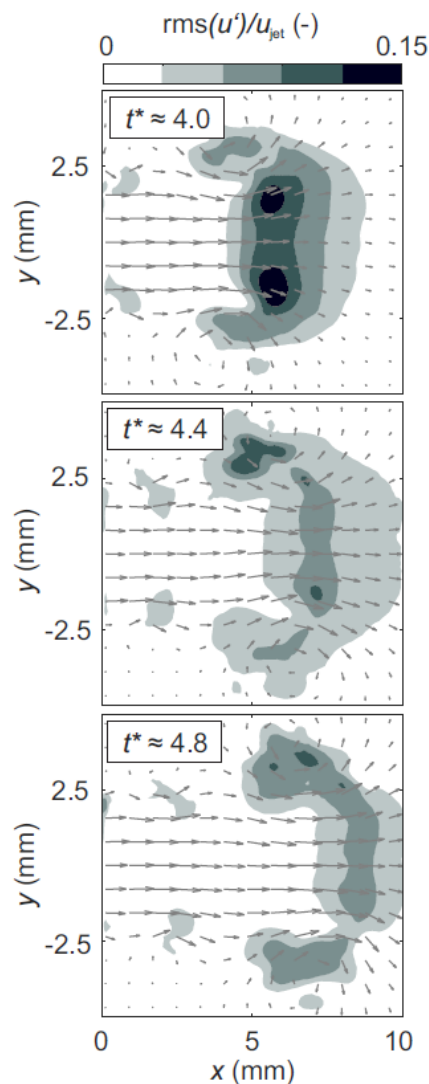


Figure 19 Time series of pulsed jet; top: phase-averaged axial velocity contour overlaid with in-plane velocity vectors, locations of vortex core highlighted by red and blue circles; bottom: probability density functions of vortex core outlet distances for the displayed phases

While in the majority of the measurement domain, velocity fluctuations are small compared to the jet exit velocity (mostly below 3%, i.e. on the order of the measurement uncertainty), 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.