

Session 9

Inference, Sensing, Inverse Modelling

- S9-A1** Sebastian Sachs — Artificial neural networks for event-based particle tracking using STELLA
- S9-A2** Mahdi Jamshidiha — Bayesian Calibration of Turbulence-Chemistry for OH* Prediction in Hydrogen Combustion
- S9-A6** Karol Bukowski — Data-driven modelling and state estimation of sloshing dynamics from sparse measurements: an experimental and numerical study
- S9-A7** Jiawei Wu — Geometry-Generalizable Flow Reconstruction via Implicit Neural Representations
- S9-P4** Alexander Stroh — Predicting 4D Mitral Regurgitation Hemodynamics from Sparse Data using Hybrid Deep Operator Networks
- S9-P5** John Agyapong Boamah — Reconstruction of Missing Velocity Fields Using Linear Modal Decomposition and Non-linear Deep Learning Approaches

Artificial neural networks for event-based particle tracking using STELLA

¹Sebastian Sachs*, ^{2,3}Steffen Jung; ^{2,3}Margret Keuper; ¹Christian Cierpka

¹Technische Universität Ilmenau, Institute of Thermodynamics and Fluid Mechanics, Ilmenau, Germany

²University of Mannheim, Data and Web Science Group, Mannheim, Germany

³Max-Planck-Institute of Informatics, Saarland Informatics Campus, Germany

Abstract

Combining event-based vision and artificial neural networks opened new perspectives in spatiotemporal perception tasks, as the asynchronous and sparse nature of the obtained event-stream naturally aligns with learning-based processing of dynamic scenes. Especially for particle tracking velocimetry (PTV) in turbulent fluid flows, the superior temporal resolution, high dynamic range, and low data rate of event-based cameras complement the low inference time and strong performance of neural networks for object detection and tracking. However, particle detection based on the event-stream poses challenges beyond common object detection tasks, as a large number of tiny particles needs to be detected simultaneously within the field of view. In this contribution, the performance of three neural network architectures for event-based particle detection is evaluated: a recurrent vision transformer, a spiking vision transformer, and a mixture-of-experts heat conduction-based detector. To provide a baseline for comparison, we introduce a modular framework for spatiotemporal event-based Lagrangian particle tracking (STELLA), which integrates direct processing of the asynchronous event-stream and processing based on dense pseudo-frame representations. The networks are trained on synthetic and experimental datasets with rich ground truth on particle position and velocity. Among the experimental datasets, a printed particle pattern mounted on a chopper wheel rotating at up to 49.5 1/s and the particle motion in a cylinder wake at Reynolds numbers up to 41600 (see Fig. 1) provide a demanding benchmark for event-based PTV. Both, the datasets and the framework STELLA are released publicly to facilitate a broader adoption of event-based vision for particle tracking. Owing to the low inference time and outstanding performance in particle detection and tracking, artificial neural networks substantially enhance the capabilities of current processing pipelines used for event-based PTV.

Keywords: event-based camera, event-based vision, particle tracking velocimetry, object tracking




Figure 1: Particle tracks detected in the wake of a cylinder by directly processing the event-stream. A main flow was applied in positive x-direction. The streamwise velocity component u_x of the particles is color-coded in the Lagrangian tracks. In the background, a pseudo-frame containing events accumulated over 4ms is shown.

Bayesian Calibration of Turbulence-Chemistry for OH* Prediction in Hydrogen Combustion

^{1,2}M. Jamshidiha*; ^{1,2}A. Procacci; ^{1,2}R. Amaduzzi; ^{1,2,3}A. Parente; ^{1,2}A. Coussement

¹ATM Laboratory, Université Libre de Bruxelles, Brussels, 1000, Belgium

²Brussels Institute for Thermal-fluid Systems and Clean Energy (BRITE), Belgium

³WEL Research Institute, Wavre, Belgium

Abstract

Reliable prediction of OH* chemiluminescence is critical for validating computational models of hydrogen-enriched combustion. This work introduces a data-driven framework that couples CFD simulations with experimental measurements to find the optimal turbulence–chemistry parameters for OH* field prediction. A four-dimensional uncertainty quantification (UQ) study was performed on a stagnation-point reverse-flow combustor, varying air preheat temperature, diluent mole fraction in fuel composition (dilution ratio), the turbulence constant C_{2e} , and the combustion mixing timescale coefficient C_{mix} . The resulting CFD database was reduced through Proper Orthogonal Decomposition (POD) and used to train a Gaussian Process Regression (GPR) surrogate. The surrogate accurately reconstructed unseen OH* distributions and provided quantified predictive uncertainty. Experimental OH* chemiluminescence images, corrected by inverse Abel transformation, were assimilated using a Bayesian inference scheme. The inferred posterior parameters yielded uncertainty-aware optimal turbulence and combustion parameters, aligning simulated and measured OH* chemiluminescence fields. This hybrid approach provides a statistically consistent path toward data-driven calibration of turbulence–chemistry interactions in modern combustors.

Keywords: Hydrogen combustion; OH* chemiluminescence; Surrogate modeling; Bayesian inference

*Corresponding Author, M. Jamshidiha mahdi.jamshidiha@ulb.be

Data-driven modelling and state estimation of sloshing dynamics from sparse measurements: an experimental and numerical study

¹K. Bukowski*; ¹T. De Maria; ¹S. Ahizi; ¹M. A. Mendez;

¹ *Department of Environmental and Applied Fluid Dynamics, Sint-Genesius-Rode,
von Karman Institute for Fluid Dynamics, BE.*

Keywords: Experimental Fluid Mechanics, Sloshing, SPH, Free-surface Measurement, Sparse Sensing.

Abstract. This study introduces a data-driven methodology for real-time reconstruction of the sloshing state from sparse free-surface measurements. Numerical simulations using Smoothed Particle Hydrodynamics (SPH) of a set of sloshing responses for the given tank geometry first provide space- and time-resolved data that characterise the free-surface deformation under sloshing excitation. Proper Orthogonal Decomposition (POD) is applied to this dataset to extract orthogonal bases of the free-surface shape, and only the dominant modes are retained. Optimal placement of wave elevation sensors is identified using a QR factorisation with pivoting algorithm applied to the POD bases. This algorithm chooses the most informative locations that maximise reconstruction accuracy for a limited number of sensors. Sparse elevation measurements are then taken at those locations. The full free-surface state is reconstructed by mapping the sparse measurements onto the POD bases at a given time instant. The methodology is validated against camera-based free-surface tracking, showing accurate reconstruction with only a limited number of sensors. This data-driven free-surface reconstruction methodology enables the development of active sloshing control applications in complex geometries and flow conditions.

1. Introduction

Sloshing is the oscillatory motion of a liquid free surface in a partially filled container under external excitation. In many engineering systems, such as spacecraft fuel tanks, marine vessels, and transport containers, sloshing generates large forces on tank walls and induces centre-of-mass displacements, affecting vehicle stability and structural integrity, especially when excitation triggers resonance. Sloshing mitigation strategies can be classified as passive or active. Passive control uses internal baffles to increase energy dissipation and modify the system's natural frequencies, whereas active control applies control actions through actuated baffles or controlled tank motion to counteract the sloshing response based on the estimated system state. In active control, performance strongly depends on reliable real-time state estimation used as input to the control law. This has been demonstrated for both model-based approaches by Konopka et al. [1] and machine-learning-based strategies by Ye et al. [2]. However, real-time free-surface measurements provide only sparse, localised data from elevation gauges and do not fully describe the free-surface state. This limitation motivates the development of methods that reconstruct the complete free surface from limited measurements in real time.

2. Methodology

The present work adopts the framework described by Manohar et al. [3] for reconstructing high-dimensional fields from sparse measurements using low-rank bases obtained from data-driven modal decomposition. The free-surface elevation field is represented using Proper Orthogonal Decomposition (POD), which separates spatial modes from their temporal coefficients :

By retaining only the first dominant modes, a low-rank approximation of the field is obtained:

This approximation captures the dominant sloshing dynamics while significantly reducing the system dimension. To identify a limited number of sensor locations such that the collected signals are maximally informative, QR factorisation with column pivoting is applied to the truncated spatial POD modes. This procedure selects optimal sampling locations tailored to the considered sloshing conditions, and the measurements are collected only at these locations. The temporal coefficients are then estimated from the sparse measurements using a least-squares projection onto the same modal basis. Finally, the full free-surface field is reconstructed in real time by combining the truncated spatial modes with their estimated temporal coefficients :

This approach allows real-time reconstruction of the full free surface from a small number of sensors.

3. Numerical Simulations

Three-dimensional sloshing simulations are performed using the SPH solver DualSPHysics developed by Domínguez et al. [4]. The free surface is extracted from the simulations and used as training data, which is decomposed using POD to obtain a reduced-order basis that captures the dominant dynamics.

4. Experimental Setup

A rectangular tank ($30 \times 45 \times 50$ cm) was mounted on the VKI Shakespeare 3-DOF sloshing table, with a maximum acceleration of 1g at 10 Hz. Sparse free-surface measurements are obtained using wave gauges placed at numerically identified optimal locations and fed into the model to reconstruct the free surface. Simultaneously, the full free surface is captured by a high-speed camera and interface tracking. The accuracy of the reconstructed free surface is assessed by comparison with free-surface tracking results.

5. Expected outcome

The framework is first tested on a canonical rectangular tank and then extended to more complex flow settings. For a rectangular tank, the sloshing motion can be described using only a few modes, allowing reconstruction with a small number of sensors. When internal baffles are added, the flow becomes more complex, requiring more modes and sensors to maintain accuracy. This work aims to evaluate reconstruction accuracy, determine the minimum number of sensors, and investigate the advantages of optimal sensor placement. The results are expected to show that real-time, data-driven reconstruction of sloshing dynamics is possible for more complex geometries. In future work, the methodology will be integrated into the Reinforcement Twinning framework for control applications.

Acknowledgements

This project is supported by the European Research Council (ERC) under the Horizon Europe programme (Starting Grant No. 101165479, RE-TWIST).

6. References

- [1] M. Konopka, F. De Rose, H. Strauch, C. Jetzschmann, N. Darkow and J. Gerstmann ‘Active slosh control and damping – Simulation and experiment’, *Acta Astronautica*, Vol. 158, pp. 89–102, 2019; <https://doi.org/10.1016/j.actaastro.2018.06.055>
- [2] M. Ye, Y. Ren, S. Zhang, H. Ma, X. Hu and O. J. Haidn ‘Sloshing suppression with a controlled elastic baffle via deep reinforcement learning and SPH simulation’, arXiv:2505.02354, 2025
- [3] K. Manohar, B. W. Brunton, J. N. Kutz and S. L. Brunton ‘Data-Driven Sparse Sensor Placement for Reconstruction: Demonstrating the Benefits of Exploiting Known Patterns’, *IEEE Control Systems Magazine*, Vol. 38, No. 3, pp. 63–86, 2018; <https://doi.org/10.1109/MCS.2018.2810460>
- [4] J. M. Domínguez, G. Fourtakas, C. Altomare, R. B. Canelas, A. Tafuni, O. García-Feal, I. Martínez-Estévez, A. Mokos, R. Vacondio, A. J. C. Crespo, B. D. Rogers, P. K. Stansby and M. Gómez-Gesteira ‘DualSPHysics: from fluid dynamics to multiphysics problems’, *Computational Particle Mechanics*, 2021; <https://doi.org/10.1007/s40571-021-00404-2>

[5] L. Schena, P. A. Marques, R. Poletti, S. Ahizi, J. Van den Berghe and M. A. Mendez ‘Reinforcement Twinning: From digital twins to model-based reinforcement learning’, Journal of Computational Science, Vol. 82, 102421, 2024; <https://doi.org/10.1016/j.jocs.2024.102421>

Geometry-Generalizable Flow Reconstruction via Implicit Neural Representations

¹Jiawei Wu; ^{1,2}Wang Han*; ^{1,2}Lijun Yang

¹*School of Astronautics, Beihang University, Beijing, China;*

²*National Key Laboratory of Aerospace Liquid Propulsion, Beihang University, Beijing, China*

Abstract

We propose a geometry-aware implicit neural representation framework for geometry-generalizable flow reconstruction from sparse observations. The framework is able to use sparse observations of the two velocity components at only 5% of the points, together with coordinates, geometric information, and flow conditions, to reconstruct complete physical states under different geometric configurations. The framework is evaluated on the FlowBench LDC_NSHT_2D_variable-Re, CYLINDERFLOW, and AIRFOIL datasets, covering both structured-grid and irregular-mesh settings as well as steady and time-dependent flows. The results show that the proposed coordinate-conditioned representation can recover the dominant flow structures and reconstruct unobserved variables across different geometrical configurations.

Keywords: Flow Reconstruction; Implicit Neural Representations; Geometry Generalization

Predicting 4D Mitral Regurgitation Hemodynamics from Sparse Data using Hybrid Deep Operator Networks

Jakob Marcel Hoffmann¹, Yosuke Hasegawa², and Alexander Stroh¹

¹Institute of Fluid Mechanics, Karlsruhe Institute of Technology, 76131 Karlsruhe, Germany

²Center for Research on Innovative Simulation Software, Institute of Industrial Science, The University of Tokyo, Japan

Abstract

Accurate quantification of mitral regurgitation (MR) is essential for clinical decisions and requires analyzing complex hemodynamics. 2D Particle Image Velocimetry (PIV) produces spatially sparse data, and transient Computational Fluid Dynamics (CFD) simulations remain computationally prohibitive for routine clinical use. This work proposes a hybrid Deep Operator Network (DeepONet) framework to reconstruct full transient volumetric (4D) hemodynamics from sparse 2D planar velocity slices and time-resolved pressure measurements. The model learns a generalized solution operator from Unsteady Reynolds-averaged Navier-Stokes (URANS) simulations and uses rapid Test-Time Adaptation (TTA) to fine-tune predictions directly on case-specific sparse observations. This approach infers plausible 3D jet trajectories in minutes. Validation on unseen synthetic data shows the model can correct in-plane flow topologies for out-of-distribution geometries. Currently, the TTA correction remains strongly localized to the supervised 2D measurement plane, causing volumetric coherence to degrade sharply away from this area and creating disjointed 3D flow structures when the pretrained baseline predictions differ significantly from the ground truth. The application of this framework to in-vitro measurements amplifies non-physical artifacts driven by inherent measurement noise and flexible boundary conditions, underscoring its susceptibility to data outside the training distribution.

1 Introduction

Mitral regurgitation (MR) causes retrograde blood flow due to inadequate valve closure. Accurate severity quantification is crucial for clinical decisions but is often underestimated by standard diagnostics (Enriquez-Sarano et al., 2009). To systematically study MR, an in-vitro simulator using Mitral Regurgitation Orifice Phantoms (MROPs) can be used to replicate leaky valves under highly controlled, physiological conditions (Karl et al., 2024; Leister et al., 2025). Analyzing the flow dynamics in a real simulator remains challenging since capturing 3D velocity fields e.g. via 2D Particle Image Velocimetry (2D2C PIV) is very time-consuming and optically restricted, while computational fluid dynamics (CFD) simulations are far too computationally intensive for routine clinical use. As a promising CFD surrogate (H. Wang et al., 2024), Deep Operator Networks (DeepONets) learn general parametric mappings for near-instant inference (Lu et al., 2021), avoiding the costly retraining required by Physics-Informed Neural Networks (Raissi et al., 2019). Because DeepONet accuracy can degrade on out-of-distribution inputs (Goswami et al., 2023), hybrid “Test-Time Adaptation” (TTA) paradigms can be utilized to rapidly fine-tune pre-trained models using sparse, case-specific observations (Liang et al., 2025; Zhu et al., 2023).

Building on this, we propose a hybrid DeepONet to reconstruct complete transient 4D hemodynamics using only planar 2D2C PIV velocity slices and pressure measurements. Pre-trained on a URANS CFD database of various MROPs, the generalized operator undergoes rapid TTA fine-tuning on the sparse target measurements. By mapping a 2D2C slice to a full 4D flow field, this framework bypasses lengthy PIV acquisitions and high CFD costs, enabling highly efficient flow analysis, which might be integrated into clinical evaluations.




Figure 1: Hemodynamic simulator geometry: (a) meshed internal fluid domain, (b) mesh slices, (c) magnified Pinhole-L orifice mesh. MROP shapes with height and width markers are shown in (d). LV = left ventricle, LA = left atrium.

2 Methodology

2.1 Synthetic Training Data

The proposed framework relies on a simulated hemodynamics database mirroring the physical in-vitro simulator (Fig. 1a). Using experimental fluid properties and transient mass flow boundary conditions, the orifice plate was modeled as a rigid wall to maintain computational tractability. Simulations were performed in Simcenter STAR-CCM+ Version 2406 (19.04.009), solving the Unsteady Reynolds-Averaged Navier-Stokes (URANS) equations with the $k-\omega$ -SST model (Menter, 1994). The domain was discretized into an optimized 400,000-cell polyhedral mesh (Fig. 1b,c), requiring approximately 60 hours of compute time per cardiac cycle on 64 cores (AMD EPYC 7662).

The dataset comprises 11 Mitral Regurgitation Orifice Phantoms (MROPs, Fig. 1d) across three geometric classes (Slot, Drop, Pinhole) at varying severities represented through different sizes (S,M,L,XL), with 300 snapshots spanning one 80 BPM cycle each. This diverse array establishes a “prior knowledge” distribution of transvalvular hemodynamics, enabling inference in unobserved regions. The oversized Drop-XL and the novel Slot-L-Bent (simulating maximum foil deformation) were excluded for validating the framework’s extrapolation performance.

2.2 DeepONet Architecture

To reconstruct 4D flow fields, a Deep Operator Network (DeepONet, Lu et al., 2021) is employed to learn a general parametric mapping between input and output functions via two parallel subnetworks (Fig. 2). The branch network input is a flattened, spatially masked 2D velocity slice (u, v) at a specific time step t , mimicking the restricted 2D2C PIV field of view. Concurrently, the trunk network ingests continuous spatiotemporal inference coordinates (x, y, z, t). Their outputs are merged via elementwise multiplication to predict flow quantities.

To improve convergence, physical quantities are non-dimensionalized ($\alpha^* = \alpha/\alpha_c$) using characteristic scales: $l_c = 0.1$ m, $t_c = 0.75$ s, $u_c = 4.988$ m/s, $p_c = 13439.05$ Pa, and $\nu_{t,c} = 3.92 \cdot 10^{-4}$ m²/s. Both networks are Multilayer Perceptrons (MLPs) featuring a high-capacity 9×640 architecture with adaptive SiLU (Swish) activations to capture highly non-linear jet dynamics (Elfwing et al., 2018; Jagtap et al., 2020).




Figure 2: Schematic of the DeepONet architecture for 3D transient flow reconstruction, adapting the branch and trunk network structure for planar PIV inputs. All inputs and outputs are non-dimensional; notation $\square^*$ is omitted for legibility.

2.3 Training Strategy

Training follows a two-step “Test-Time Adaptation” (TTA) paradigm using the AdamW optimizer (Loshchilov & Hutter, 2019). First, pre-training optimizes the network on the full CFD dataset to learn a generalized operator ($\mathcal{L}_{\text{data}}^{\text{train}}$). This phase requires 50,000 steps and approximately 12 hours of computation (RTX 6000 Ada GPU). Second, hybrid fine-tuning rapidly adapts the pre-trained weights to a specific unseen case. This step utilizes sparse 2D PIV slices and time-resolved pressure measurements in the LA and LV ($\mathcal{L}_{\text{data}}^{\text{val}}$), alongside the ongoing training distribution. This highly efficient calibration requires only 400 additional steps (taking from 10 minutes to 1.5 hours).

The data loss formulation utilizes the Sum of Squared Errors (SSE) across velocity components (u, v, w), static pressure (p), and kinematic eddy viscosity (ν_t), balanced dynamically via Neural Tangent Kernel (NTK) weighting (S. Wang et al., 2022). Outputting ν_t directly closes the RANS equations without requiring full transport PDEs for k and ω . Although Physics-Informed residual losses ($\mathcal{L}_{\text{PDE}}$) were tested, they were ultimately omitted; purely data-driven fine-tuning proved more robust, whereas PDE constraints induced over-smoothing and high computational costs in the complex 3D domain.

3 Results

3.1 Validation on Synthetic Data

The considered hemodynamics is characterized by a high-velocity systolic regurgitant jet entering the left atrium (LA), followed by a diastolic backflow into the left ventricle (LV). While this diastolic jet is a non-physiological remnant of the fixed MROP experimental design, it provides additional insight for assessing the model’s robustness. To assess the framework’s extrapolation capabilities, the DeepONet was evaluated on the two unseen synthetic cases: Drop-XL and Slot-L-Bent. Overall, the reconstructed 4D solutions demonstrate high fidelity.

As visualized for the systolic phase of the Drop-XL case in Fig. 3, the model accurately captures the regurgitant jet’s magnitude and spatial extent. Crucially, the prediction respects the complex geometric boundaries; the jet interacts naturally with the atrial walls, rolling off the geometry rather than penetrating it, which results in low error magnitudes along the surface.




Figure 3: (a) Reference CFD velocity magnitude at $z^* = 0$ and $t^* = 0.25$ for the Drop-XL orifice. (b) DeepONet prediction. (c) Signed non-linear and (d) absolute error magnitude with streamline comparison.

This robustness to out-of-distribution inputs was further confirmed by the Slot-L-Bent validation case. Here, the framework successfully inferred the onset of the jet despite a physical coordinate shift of the orifice opening caused by the introduction of a curved foil geometry, demonstrating the operator's capacity to adapt to significant geometric variations.

Despite these strong general results, the data-driven mapping between the sparse 2D input and the dense 3D output exhibits specific structural limitations. Because the exact timing of jet formation varies across different MROPs, the pre-trained baseline model can inaccurately predict the transient state, often positioning the jet pulse temporally ahead of the ground truth. Test-Time Adaptation (TTA) effectively rectifies this by retracting the jet to match the sparse observations. However, these adjustments remain highly spatially localized. The correction at the supervised $z^* = 0$ center plane does not propagate uniformly across the full z -width of the domain. Consequently, unconstrained outer fluid layers remain in their original, advanced temporal positions, sometimes creating a disjointed 3D topology.

Furthermore, extrapolating to the larger Drop-XL orifice introduces distinct underprediction streaks in the unconstrained y -direction and LV regions, where the sparse supervision provides no direct correction. Finally, an analysis of the Normalized Root Mean Square Error (NRMSE) calculated per snapshot over the cardiac cycle reveals that the total quantitative improvement yielded by TTA is small and mostly attributed to the specific transient phases where higher velocity magnitudes occupy the supervised 2D measurement window.

3.2 Application to Experimental Data

The final framework was applied to four real-world in-vitro measurements (Pinhole-L, Slot-L, Drop-XL, and EccJet), two of which represent geometries included in the CFD training dataset. Moving beyond synthetic validation introduces significant constraints: the absence of a full 3D volumetric ground truth, temporal sparsity (fewer than 50 snapshots compared to 300 in CFD), noisy input signals, and a visible domain misalignment caused by the flexible orifice foil used in the physical experiment versus the rigid walls modeled in the CFD training distribution.




Figure 4: (a) Reference PIV velocity magnitude at $z^* = -0.02$ and $t^* = 0.067$ for the EccJet orifice. (b) DeepONet prediction. (c) Absolute error magnitude with streamline comparison.

The experimental applications broadly mirror the capabilities and limitations observed during the synthetic validation, capturing the general flow topologies. However, applying the framework to real data intensifies the previously identified sources of error and introduces more erratic, non-physical artifacts, particularly speckled velocity structures in the unconstrained areas. This degradation is likely driven by the combined effects of measurement noise, the aforementioned domain misalignment, and the database’s inability to perfectly capture the complex reality of the driving input, such as the experimental pump reaching its power limit in the DropXL case, which generated unfamiliar boundary conditions.

Due to the lack of dense ground truth, the physical consistency of the experimental reconstructions was evaluated using spatiotemporal L1 integral norms of the URANS equations. Despite the higher prevalence of visually identifiable issues, the calculated PDE residuals exhibit similar magnitude with those of the CFD validation cases, indicating that the model maintains a baseline level of physical consistency even when evaluated for experimental input.

The EccJet configuration serves as a rigorous final stress test due to its highly eccentric jet angle, completely absent in the training dataset. Following TTA, the hybrid model successfully reorients the jet direction to match the eccentric angle observed in the PIV data. Crucially, this dataset provides time-resolved PIV measurements for adjacent z -planes, enabling a quantitative assessment of the volumetric reconstruction. As shown in Fig. 4, the model successfully captures the general features near the supervised center plane ($z = -2$ mm), achieving an NRMSE of 24.6%. However, prediction accuracy degrades sharply as the distance from the supervised slice increases; at a deeper plane ($z = -6$ mm), the structure loses definition, and the NRMSE more than doubles to 52.6%, a trend of increasing error and structural degradation consistently observed across all other available timesteps. Additionally, traces of the training distribution persist in the prediction, visible as falsely elevated velocities along the right atrial wall. These findings indicate that while the model successfully repositions the jet based on planar data, the single-plane TTA strategy is currently insufficient to ensure reliable volumetric coherence across the full 3D domain.

4 Conclusion

This work proposed a hybrid Deep Operator Network (DeepONet) to reconstruct 4D hemodynamics from sparse 2D planar velocities and pressure measurements. The framework infers plausible full-field flows in minutes, bypassing prohibitive CFD computational costs. The Test-Time Adaptation (TTA) strategy effectively corrects in-plane topologies for out-of-distribution geometries, even reorienting a highly eccentric trajectory.

However, critical limitations remain. The TTA correction is strongly localized to the 2D

measurement plane and fails to propagate uniformly across the domain's depth. This causes disjointed 3D structures when the baseline prediction deviates significantly from the true state, and errors from geometric extrapolation. Additionally, the purely data-driven framework is sensitive to the domain shift between idealized, rigid-wall CFD training data and noisy, flexible-boundary experimental signals, increasing non-physical artifacts.

Future research should explore (orthogonal) multi-plane supervision, which could resolve the lateral propagation issue. Temporal consistency can likely be improved by replacing the snapshot-based approach with e.g. Temporal Convolutional Networks (TCNs). Furthermore, robustness against the simulation-to-reality gap can be enhanced by expanding the CFD database, utilizing noise-based data augmentation, and re-exploring physics-based losses. In summary, while the current approach provides a promising pathway for reduced-cost hemodynamic modeling, further development is required to achieve coherent 4D reconstructions.

5 Acknowledgments

This work was supported by the JSPS invitational fellowship for researchers in Japan, ID S25051, the UTokyo-IIS internship support program, and the DAAD-PROMOS scholarship.

References

- Elfwing, S., Uchibe, E., & Doya, K. (2018). Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural Networks*, 107, 3–11. <https://doi.org/10.1016/j.neunet.2017.12.012>
- Enriquez-Sarano, M., Akins, C. W., & Vahanian, A. (2009). Mitral regurgitation. *The Lancet*, 373(9672), 1382–1394. [https://doi.org/10.1016/S0140-6736\(09\)60692-9](https://doi.org/10.1016/S0140-6736(09)60692-9)
- Goswami, S., Bora, A., Yu, Y., & Karniadakis, G. E. (2023). Physics-Informed Deep Neural Operator Networks. In T. Rabczuk & K.-J. Bathe (Eds.), *Machine Learning in Modeling and Simulation: Methods and Applications* (pp. 219–254). Springer International Publishing. https://doi.org/10.1007/978-3-031-36644-4_6
- Jagtap, A. D., Kawaguchi, K., & Karniadakis, G. E. (2020). Adaptive activation functions accelerate convergence in deep and physics-informed neural networks. *Journal of Computational Physics*, 404, 109136. <https://doi.org/10.1016/j.jcp.2019.109136>
- Karl, R., et al. (2024). An ex-vivo and in-vitro dynamic simulator for surgical and transcatheter mitral valve interventions. *International Journal of Computer Assisted Radiology and Surgery*, 19(3), 411–421. <https://doi.org/10.1007/s11548-023-03036-4>
- Leister, R., et al. (2025). Investigating the shortcomings of the flow convergence method for quantification of mitral regurgitation in a pulsatile in-vitro environment and with computational fluid dynamics. *Cardiovascular Engineering and Technology*. <https://doi.org/10.1007/s13239-024-00763-w>
- Liang, J., He, R., & Tan, T. (2025). A comprehensive survey on test-time adaptation under distribution shifts. *International Journal of Computer Vision*, 133(1), 31–64. <https://doi.org/10.1007/s11263-024-02181-w>
- Loshchilov, I., & Hutter, F. (2019). Decoupled weight decay regularization. *International Conference on Learning Representations (ICLR)*. <https://arxiv.org/abs/1711.05101>
- Lu, L., Jin, P., Pang, G., Zhang, Z., & Karniadakis, G. E. (2021). Learning nonlinear operators via deepnet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3, 218–229. <https://doi.org/10.1038/s42256-021-00302-5>
- Menter, F. R. (1994). Two-equation eddy-viscosity turbulence models for engineering applications. *AIAA Journal*, 32(8), 1598–1605. <https://doi.org/10.2514/3.12149>
- Raissi, M., Perdikaris, P., & 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. <https://doi.org/10.1016/j.jcp.2018.10.045>
- Wang, H., et al. (2024). Recent advances on machine learning for computational fluid dynamics: A survey. <https://doi.org/10.48550/arXiv.2408.12171>
- Wang, S., Yu, X., & Perdikaris, P. (2022). When and why PINNs fail to train: A neural tangent kernel perspective. *Journal of Computational Physics*, 449, 110768. <https://doi.org/10.1016/j.jcp.2021.110768>
- Zhu, M., et al. (2023). Reliable extrapolation of deep neural operators informed by physics or sparse observations. *Computer Methods in Applied Mechanics and Engineering*, 412, 116064. <https://doi.org/10.1016/j.cma.2023.116064>

Reconstruction of Missing Velocity Fields Using Linear Modal Decomposition and Non-linear Deep Learning Approaches

¹J. A. Boamah*; ^{1,2}R. M. Barron; ^{1,3}R. Balachandar;

¹*Department of Mechanical, Automotive and Materials Engineering, University of Windsor, Windsor, Canada;*

²*Department of Mathematics and Statistics, University of Windsor, Canada;*

³*Department of Civil and Environmental Engineering, University of Windsor, Canada*

Abstract

Reconstruction of incomplete velocity fields remains a significant challenge in experimental fluid mechanics due to missing or occluded regions in particle image velocimetry (PIV) measurements. This study evaluates linear and nonlinear supervised learning techniques for reconstructing missing wake regions behind a horizontal circular cylinder, positioned near a bed with a clearance of 25% of the diameter at a Reynolds number of 8000. An extended proper orthogonal decomposition (EPOD) model and a U-Net convolutional neural network are evaluated using experimentally acquired PIV datasets with synthetic masking corresponding to 5% and 10% streamwise gaps. The influence of POD-based denoising during training is investigated by comparing models trained using actual and denoised PIV fields. The results show that the U-Net models consistently outperform EPOD for all gaps. Although all methods recover the dominant large-scale wake dynamics, EPOD exhibits substantial degradation in the energy content at low and intermediate wavenumber and intermittent vortical motions. Models trained using POD-denoised fields exhibit smoother training behaviour, improved spectral consistency, and closer agreement with the reference vorticity statistics, indicating enhanced robustness to experimental noise. In contrast, models trained using actual experimental data achieve lower pointwise reconstruction error, highlighting a trade-off between numerical accuracy and preservation of turbulence physics.

Keywords: Reconstruction, EPOD, U-Net, PIV

1. Introduction

In experimental fluid mechanics, particle image velocimetry (PIV) is widely used technique to obtain planar velocity measurements [1]. However, experimental limitations such as laser reflections, inadequate particle seeding, and geometric obstructions often produce missing or corrupted regions within the measured flow field. Reconstruction of incomplete velocity fields remains a significant problem, and techniques for recovering incomplete velocity information are essential for reliable flow characterization and analysis. Conventional reconstruction methods have largely been based on interpolation and linear estimation techniques. Proper orthogonal decomposition (POD) was introduced as a reduced-order framework capable of representing complex turbulent flows using a limited number of dominant energetic modes [2] [3]. Based on this, Everson and Sirovich [4] proposed the gappy POD approach, which reconstructs incomplete datasets by projecting sparse measurements onto POD basis derived from complete flow snapshots. Later, extended POD (EPOD) incorporated cross-correlation relation between measured and unmeasured quantities to improve reconstruction performance [5]. Although POD-based techniques are computationally efficient and physically interpretable, their linear formulation limits their ability to capture strongly nonlinear and multi-scale flow structures.

Machine learning approaches have recently been applied to reconstruct incomplete PIV measurements and flow-field inpainting problems. Early studies used supervised methods such as multilayer perceptrons and support vector regression to estimate missing velocity data with improved accuracy over conventional interpolation techniques [6] [7]. More recently, convolutional neural networks (CNNs) and encoder-decoder architectures, including U-Net models, have demonstrated strong capability in recovering nonlinear flow structures and reconstructing missing velocity regions in turbulent flows [8]. In the present study, linear and nonlinear reconstruction methods are compared under supervised learning for recovering missing velocity regions in the wake of a horizontal circular cylinder. An EPOD-based model and a U-Net deep learning

*Corresponding Author, J. A. Boamah boamahj@uwindsor.ca

architecture are evaluated using experimentally obtained PIV datasets with synthetic masking conditions. Additionally, POD-based denoising is employed during training to investigate its influence on reconstruction performance under realistic experimental conditions.

2. Methodology

The present study investigates reconstruction of missing regions in planar particle image velocimetry (PIV) velocity fields using both linear modal decomposition and nonlinear deep learning approaches. The datasets were obtained from experimental measurements of the near wake of a circular cylinder located near a bed, as shown in figure 1a, with a clearance ratio of $d = 8000$ [9]. A total of 4,000 instantaneous velocity fields is used, where 3,200 samples are allocated for training and validation and 800 samples are reserved for testing. Each velocity field consists of streamwise and vertical velocity components (u, v) defined on a structured 150×200 grid. The data are normalized using z-score normalization, and synthetic masks corresponding to 5% and 10% streamwise gaps, as shown in figure 1b, are imposed to emulate missing measurement regions. To investigate the effect of noise on the supervised reconstruction, POD-based denoising is applied only to the training dataset, while all model evaluations are performed using the actual PIV measurements.

For linear reconstruction, extended proper orthogonal decomposition (EPOD) is applied following the formulation of Borée [5] and Li et al. [10]. POD is first applied to the measured region to obtain dominant spatial modes and corresponding temporal coefficients through singular value decomposition (SVD) using 95% of the turbulent kinetic energy. Cross-correlations between the modal coefficients and the missing region are then used to construct EPOD modes, enabling estimation of the missing velocity field from the available measurements. The reconstructed field is subsequently recovered by superposition of the retained modes and the mean flow field.




Figure 1 Schematic of (a) the flow features, (b) the masked region defined for training, and (c) the network used for the velocity field reconstruction.

A U-Net-based convolutional neural network [11] is used as the nonlinear reconstruction model as shown in figure 1c. The network input consists of the normalized velocity components together with a binary mask identifying missing regions, while the output corresponds to the reconstructed velocity field. The architecture follows a symmetric encoder-decoder configuration with skip connections to preserve both large-scale and fine-scale flow structures. The encoder progressively extracts hierarchical flow features through convolution and downsampling operations, whereas the decoder restores spatial resolution through upsampling and feature fusion through the skip connections. The network is trained in a supervised framework using synthetically masked velocity fields, with optimization performed using the Adam optimizer and mean squared error (MSE) loss. Five-fold cross-validation, learning-rate scheduling, and early stopping are used to improve robustness and reduce overfitting. Models herein trained using PIV data

*Corresponding Author, J. A. Boamah boamahj@uwindsor.ca

are referred to as UNet-PIV, whereas models trained using POD-denoised datasets reconstructed with 95% energy retention are referred to as UNet-dPIV.

3. Results and Discussion

Figure 2 presents the instantaneous reconstructed velocity fields and corresponding relative error distributions for the EPOD, UNet-PIV, and UNet-dPIV models for a gap size of 10%. All models successfully recover the overall wake topology within the missing region. However, clear differences emerge in the reconstruction of localized flow structures. The EPOD model exhibits larger errors in the near-bed jet region, where strong velocity gradients and complex shear-layer interactions dominate the flow. Both U-Net-based models produce smoother and more spatially coherent reconstructions. These observations are quantified in table 1 through the mean absolute error (MAE) computed over all test snapshots within the masked region.

For both gap sizes, the U-Net models consistently outperform EPOD for reconstruction of the streamwise (u) and vertical (v) velocity components. At the 5% gap size, all models reconstruct the dominant streamwise motion with relatively low error, whereas the v -component shows substantially higher MAE values due to the stronger intermittency and smaller spatial scales associated with transverse wake motion. As the gap size increases from 5% to 10%, the reconstruction error increases for all methods, demonstrating the greater difficulty of recovering flow information over wider missing regions. The degradation is most significant for EPOD, particularly for the v -component where the MAE increases from 19.7% to 43.3%, highlighting the limitation of linear modal reconstruction in representing highly nonlinear wake dynamics. Although the difference between UNet-PIV and UNet-dPIV remains modest, the denoised training data consistently produce smoother reconstructions and slightly lower errors, suggesting that POD filtering improves the stability of the learned flow representation.




Figure 2 Instantaneous streamwise (u) and vertical (v) velocity fields showing the reconstructed and ground-truth data for 10% across all models.

Table 1 Relative mean absolute error (MAE) for EPOD and PIV reconstructions within the missing region for 5% and 10% gap sizes.

MAE (5%) - %			MAE (10%) - %		
EPOD	UNet-PIV	UNet-dPIV	EPOD	UNet-PIV	UNet-dPIV
7.3	4.4	4.7	15.3	10.2	10.5
19.7	12.8	13.6	43.3	32.6	33.8

The trends observed in the instantaneous fields and MAE analysis are further reflected in the spectral behaviour shown in figure 3a. The spectral ratio, defined as the reconstructed streamwise energy relative to the ground truth, provides a scale-resolved assessment of reconstruction accuracy within the gap region. At low wavenumbers, all models remain close to unity, indicating that the dominant large-scale wake structures are accurately recovered even with missing data. Differences become increasingly pronounced at intermediate and high wavenumbers, where the reconstruction of small-scale turbulent motions becomes more challenging. EPOD progressively deviates from unity, consistent with the elevated MAE values and the visible smoothing observed in figure 2. The UNet-PIV model maintains improved agreement over a broader spectral range but exhibits small oscillations associated with measurement noise retained in the PIV training data. UNet-dPIV remains closest to unity across the widest range of wavenumbers, demonstrating that POD-based denoising enables the network to better preserve physically relevant flow structures while reducing noise contamination. Increasing the gap size to 10% systematically degrades the high-wavenumber reconstruction for all models, with the largest deterioration observed for EPOD.




Figure 3 (a) Comparison of gap-averaged energy spectra and (b) Vorticity PDFs between the reconstructed models and the ground truth at $x/d=1.5$ and $y/d=1.1$ for all missing gaps, compared across all models

To further examine the reconstruction of small-scale turbulent dynamics, figure 3b compares the vorticity probability density functions (PDFs) at . The ground-truth distribution exhibits the broad heavy-tailed behaviour characteristic of intermittent shear-layer vortices. Consistent with the spectral results, EPOD shows the largest deviation from the reference distribution, particularly in the tails, indicating underprediction of intense vorticity fluctuations. Both U-Net models provide closer agreement with the ground truth, confirming their improved capability in reconstructing intermittent flow structures. Among them, UNet-dPIV yields the best agreement in both the peak and tail regions, further supporting the conclusion that denoised training data improve recovery of localized vortical motions and high-frequency flow content. Overall, the instantaneous, statistical, spectral, and vorticity analyses consistently demonstrate

*Corresponding Author, J. A. Boamah boamahj@uwindsor.ca

that nonlinear deep learning approaches outperform linear EPOD reconstruction for incomplete PIV fields, while POD-based denoising provides additional robustness under realistic experimental conditions.

4. Conclusions

This study compares linear EPOD and nonlinear U-Net approaches for reconstructing missing velocity regions in experimental PIV measurements of the wake flow behind a horizontal circular cylinder with a bed clearance of 0.25 at . The results demonstrate that both approaches can recover the dominant wake structures. However, the nonlinear U-Net models consistently provide lower reconstruction errors and improved recovery of localized turbulent features. The limitations of EPOD become increasingly pronounced as the gap size increases, particularly for the vertical velocity component and high-wavenumber flow content. Spectral and vorticity analyses further show that the U-Net models better preserve large- and intermediate-scale turbulent dynamics and intermittent vortical structures. Although the improvement between UNet-PIV and UNet-dPIV is modest, POD-based denoising during training produces smoother reconstructions and improved spectral stability, indicating enhanced robustness to experimental noise. The present results demonstrate the potential of supervised deep learning approaches for reliable reconstruction of incomplete experimental velocity fields and provide a framework for future applications in complex turbulent flows.

5. References

- [1] R. J. Adrian (1991). Particle imaging techniques for experimental fluid mechanics. *Annual Review of Fluid Mechanics*, 23:261–304.
- [2] J. L. Lumley (1967). The structure of inhomogeneous turbulent flows. In: A. M. Yaglom and V. I. Tatarski, editors. *Atmospheric Turbulence and Radio Wave Propagation*, pages 166–178. Nauka, Moscow.
- [3] G. Berkooz, P. Holmes, and J. L. Lumley (1993). The proper orthogonal decomposition in the analysis of turbulent flows. *Annual Review of Fluid Mechanics*, 25:539–575.
- [4] R. Everson and L. Sirovich (1995). Karhunen–Loève procedure for gappy data. *Journal of the Optical Society of America A*, 12(8):1657–1664.
- [5] J. Borée (2003). Extended proper orthogonal decomposition: A tool to analyse correlated events in turbulent flows. *Experiments in Fluids*, 35(2):188–192.
- [6] G. Akbari and N. Montazerin (2022). Reconstruction of particle image velocimetry data using flow-based features and validation index: A machine learning approach. *Measurement Science and Technology*, 33(1):015203.
- [7] M. H. Aksoy, I. Goktepe, M. Ispir, and A. Cakan (2023). Reconstruction of flow field with missing experimental data of a circular cylinder via machine learning algorithm. *Physics of Fluids*, 35(11):114110.
- [8] C. Li, Y. Wan, B. Liang, Q. Zhang, M. Zhao, Y. Liu, and P. Yuan (2025). Reconstruction of incomplete velocity field and prediction of pressure field of propeller wake based on deep learning. *Physics of Fluids*, 37(7):075177.
- [9] S. Bista, G. Nasif, and R. Balachandar (2025). Interference effects on the wake dynamics of a horizontal circular cylinder close to a wall. *Physics of Fluids*, 37(10):105159.
- [10] T. Li, M. Buzzicotti, L. Biferale, F. Bonaccorso, S. Chen, and M. Wan (2023). Multi-scale reconstruction of turbulent rotating flows with proper orthogonal decomposition and generative adversarial networks. *Journal of Fluid Mechanics*, 971:A3.
- [11] O. Ronneberger, P. Fischer, and T. Brox (2015). U-Net: Convolutional networks for biomedical image segmentation. *Lecture Notes in Computer Science*, 9351:234–241.

*Corresponding Author, J. A. Boamah boamahj@uwindsor.ca