Chapter 3
Neural Operators for Accelerating Flow Field Predictions in an Empty Duct

Alan Xavier and Ludovic Renson

Abstract Scientific machine learning leverages modern machine learning (ML) methods to accelerate the analysis of dynamical systems governed by ordinary and partial differential equations. In fluid mechanics, these techniques offer the potential to approximate complex, nonlinear flow behaviour at a fraction of the cost of conventional computational fluid dynamics (CFD). While most prior approaches have relied on neural networks trained to learn mappings between finite-dimensional vector spaces, fluid flows are inherently continuous and infinite-dimensional. Neural operators (NOs) address this challenge by learning mappings directly between function spaces, enabling generalization across a wide range of boundary conditions and operating regimes.

In this work, we develop a neural operator framework based on the Deep Operator Network (DeepONet) to predict unsteady flow fields in a simplified aeroacoustic configuration: an empty duct with an outlet pressure boundary condition that includes a sinusoidal fluctuation. This fluctuating outlet pressure represents upstream-propagating acoustic waves, a canonical test problem for aeroelastic and turbomachinery flow analyses. High-fidelity CFD simulations are performed across a range of forcing frequencies. From these data, our DeepONet model — augmented with recurrent layers to capture temporal evolution — learns to map the transient response of the first oscillation cycle to the steady response of the final cycle.

This approach bypasses the need to simulate every cycle of the unsteady flow, reducing computational costs by an order of magnitude while retaining accuracy in the predicted density, pressure and velocity fields. The resulting operator can provide rapid flow field predictions, demonstrating a pathway for embedding scientific machine learning into turbomachinery design workflows where repeated evaluations of unsteady flow solutions are required.

Keywords Machine Learning · Neural Operators · Deep Learning · Aerodynamics · Turbomachines

Introduction

Computational fluid dynamics (CFD) is indispensable in the design and analysis of turbomachinery, where accurate predictions of unsteady flows, blade–row interactions, and aeroelastic phenomena are essential for both performance and safety. However, the fidelity required to capture the strongly coupled flow physics in rotating machinery comes at a significant computational cost. High-resolution unsteady simulations, such as unsteady Reynolds-averaged Navier–Stokes (URANS) or large eddy simulations (LES), may require thousands of CPU hours for a single operating condition. As a result, comprehensive parametric sweeps, real-time aeroelastic stability assessments, and large-scale uncertainty quantification remain infeasible within typical industrial design cycles.

Neural operators (NOs) offer a promising route to overcome these limitations by directly learning the solution operators of the governing partial differential equations. Architectures such as Fourier Neural Operators (FNOs) [1] and Deep Operator Networks (DeepONets) [2] can approximate flow fields across varying boundary conditions, blade geometries, and operating regimes. Once trained, these models can deliver high-fidelity flow predictions at a fraction of the cost of conventional solvers, enabling rapid evaluation of unsteady aerodynamic loads and modal responses. In turbomachinery applications, this

Alan Xavier · Ludovic Renson
Department of Mechanical Engineering
Imperial College London, South Kensington Campus, London SW7 2AZ, UK
e-mail: aax23@ic.ac.uk; l.renson@imperial.ac.uk

Kristin B. Zimmerman, PhD (eds.), *Data Science in Engineering, 2026, Vol. 9*, Conference Proceedings of the Society for Experimental Mechanics Series, https://doi.org/10.1007/978-8-743-81448-1_3

capability opens the door to accelerating flutter analysis, forced-response prediction, and multi-point performance optimization.

A key advantage of NOs is their ability to generalize across parametric variations without requiring retraining for each new case. This property makes them particularly well-suited for design space exploration and probabilistic risk assessment during early-stage engine development. By embedding NO-based surrogates into CFD and aeroelastic workflows, the industry can achieve both dramatic computational acceleration and enhanced predictive capability—bridging the gap between the demands of high-fidelity simulation and the real-time decision-making needs of modern propulsion system design.

Methods

We consider the prediction of unsteady flow fields in an empty duct of length L and constant cross-section, governed by the unsteady compressible Navier–Stokes equations. Different operating conditions are simulated by varying the outlet boundary condition, where a sinusoidally varying pressure wave of amplitude 300 kPa and frequency ranging from 800 Hz to 1200 Hz (in increments of 25 Hz) is applied. The inlet is modeled at atmospheric conditions to allow upstream propagation of the pressure wave. Inviscid wall conditions are enforced on the top and bottom surfaces. This canonical setup emulates aeroacoustic and aeroelastic test problems, serving as a baseline for validating operator-learning methods in unsteady turbomachinery flows.

Let $\mathbf{u}(x, y, z, t)$ denote the flow state, consisting of velocities (U_x, U_y, U_z), pressure p, and density ρ, at spatial coordinates (x, y, z) and time t. Our objective is to approximate the solution operator

$$\mathcal{G} : \mathcal{I} \mapsto \mathbf{u}(x, y, z, t),$$

which maps the transient flow field over the first oscillation period to the flow field of the final period. Instead of advancing the governing PDEs numerically, we approximate $\mathcal{G}$ using a Deep Operator Network (DeepONet) with recurrent dynamics in the branch net to capture temporal correlations.

The DeepONet framework is decomposed into two subnetworks:

1. **Branch Network (LSTM-based):** Inspired by He et al. [3, 4], the branch network encodes the temporal evolution of the inflow conditions sampled over discrete time steps during the first oscillation period. A stack of Long Short-Term Memory (LSTM) layers captures both periodicity and nonlinear dependencies of the propagating wave. This produces a latent embedding b of dimension $H \times T$, where H is the hidden dimension and T is the number of timesteps per oscillation.

2. **Trunk Network (Fully Connected):** The trunk network receives spatial coordinates (x, y, z) as input and generates a latent embedding t of dimension $N \times H \times C$, where N is the number of mesh nodes and C is the number of flow states.

3. **Operator Construction:** The operator $\mathcal{G}$ is formed by bilinear interaction of the branch and trunk embeddings:

$$\mathcal{G} = \sum_{h=1}^{H} b_{h,t}\, t_{n,h,c},$$

where $n \in [1, N]$, $h \in [1, H]$, and $c \in [1, C]$. This formulation ensures that the LSTM branch encodes inlet-driven temporal dynamics, while the trunk network captures spatial response modes of the duct.

Training Strategy

High-fidelity CFD data from the first oscillation period ($0 \leq t \leq T = 1/f$) is used for supervised training. The duct response over this window contains both the primary wave propagation and any nonlinear distortions. Since the flow is uniform in the z-direction, only samples in the xy-plane are considered for training.

To reduce input dimensionality, we sample node locations along the top wall for the branch network, as no fluctuations are observed in the y-direction. Furthermore, instead of using all five primitive variables as inputs, we compute the Mach number M, defined as

$$M = \frac{V}{a}, \quad a = \sqrt{\gamma \frac{p}{\rho}}, \quad V = \sqrt{U_x^2 + U_y^2 + U_z^2},$$

which compactly encodes the flow state. Inlet and outlet variables are normalized to $[-1, 1]$ using min–max scaling to eliminate scale discrepancies between features.

The model is trained to minimize the mean squared error (MSE) between predicted and reference CFD solutions across spatial and temporal dimensions for each flow variable:

$$\mathcal{L} = \sum_{c=1}^{C} \mathcal{L}_c, \quad c \in \{\rho, U_x, U_y, U_z, p\},$$

where each component loss is defined as

$$\mathcal{L}_c = \frac{1}{N_x N_y N_z N_t} \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} \sum_{k=1}^{N_z} \sum_{l=1}^{N_t} \left\| u(x_i, y_j, z_k, t_l) - \hat{u}(x_i, y_j, z_k, t_l) \right\|_2^2.$$

After training, the DeepONet–LSTM model is deployed to predict the flow field over the final oscillation period ($(S-1)T < t \leq ST$), for S total periods.

Results

The model was trained using mini-batches of forcing sequences covering multiple operating conditions. Optimization was performed with the AdamW optimizer [5], employing a learning rate of 1×10^{-4} and weight decay regularization. Training was conducted for 5000 epochs.

We evaluated two model variants:

1. **CFD mesh model** — trained directly on the unstructured CFD solver mesh.

2. **Structured mesh model** — trained on a uniform grid of 200×20 points in the axial and vertical directions, respectively. This structured mesh contains approximately four times fewer nodes than the unstructured CFD mesh, reducing computational expense.

Figure 1 compares the DeepONet–LSTM predictions with the reference CFD solutions for the final timestep of the last oscillation period in the duct, across three flow variables: density (ρ), axial velocity (U_x), and pressure (p).

- **Overall prediction quality:** The predicted fields (top row) closely reproduce the structure and periodicity of the CFD reference solutions (middle row). The spatial oscillations induced by the sinusoidal outlet pressure condition are captured with high fidelity across the duct length.

- **Density field (ρ):** The DeepONet model reconstructs the wave pattern with strong agreement in both wavelength and amplitude. The difference field (bottom row) shows only small discrepancies, localized near the outlet boundary (left) where the gradients are steepest.

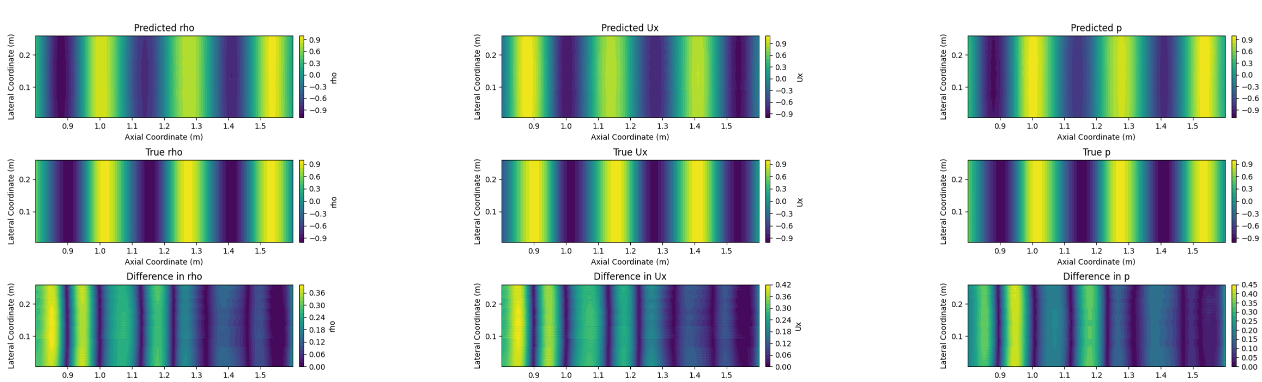

Fig. 1 Predicted flow fields for the final timestep of the last oscillation period on the **unstructured CFD mesh**. Results are shown for (a) density ρ, (b) axial velocity U_x, and (c) pressure p. For each variable, the top row shows the DeepONet–LSTM prediction, the middle row shows the reference CFD solution, and the bottom row shows the absolute difference field.

- **Axial velocity (U_x):** Predictions align well with the CFD solution, especially near the inlet (right). Minor deviations appear closer to the outlet (left), suggesting sensitivity of the network to boundary reflections or phase shifts in the velocity response.
- **Pressure field (p):** The predicted pressure oscillations match the reference solution in both phase and amplitude. Errors are again localized near the downstream boundary (left), but their magnitude remains small relative to the overall oscillation amplitude.
- **Error characteristics:** Across all three fields, the difference plots confirm that the DeepONet–LSTM model captures the dominant wave dynamics while underpredicting fine-scale variations near the outlet (left). The overall error magnitudes (≤ 0.45 in normalized units) remain modest, showing that the model generalizes well across the spatial domain.

In short, the DeepONet–LSTM framework accurately predicts the periodic wave propagation in the empty duct test case. It reproduces the dominant spatial oscillations in density, velocity, and pressure, with only minor boundary-localized discrepancies.

Figure 2 presents the DeepONet–LSTM predictions against the reference CFD solutions for the final timestep of the last oscillation period in the duct, this time on a **structured mesh (200 × 20 grid)**. As in Figure 1, results are shown for density (ρ), axial velocity (U_x), and pressure (p).

- **Overall prediction quality:** The model predictions (top row) reproduce the CFD fields (middle row) with high accuracy, and the error magnitudes (bottom row) are noticeably lower than those obtained with the unstructured mesh in Figure 1. The structured grid appears to aid the network in learning smoother and more consistent spatial patterns.
- **Density field (ρ):** The predicted density field closely follows the CFD solution across the duct. The difference field shows very small discrepancies (< 0.3 in normalized units), primarily confined to the downstream boundary (left). Compared to the unstructured mesh case, the error amplitude is reduced by roughly 25%.
- **Axial velocity (U_x):** Predictions match the CFD reference almost exactly along the duct axis. The absolute error field is smooth and minimal, indicating that the structured mesh improved the model's ability to capture phase alignment and amplitude of the oscillatory velocity field.
- **Pressure field (p):** The sinusoidal pressure oscillations are very well captured in both wavelength and amplitude. The difference field remains small and uniform across most of the duct, with maximum discrepancies near the outlet (left) but still significantly reduced compared to Figure 1.
- **Error characteristics:** The errors are uniformly small across the domain, with the largest magnitudes reduced by roughly half of those in the unstructured mesh case. The structured discretization provides cleaner spatial inputs, which likely allows the DeepONet–LSTM model to better capture the propagating wave dynamics.

In summary, the structured mesh results (Figure 2) show a clear improvement over the unstructured mesh case (Figure 1). The DeepONet–LSTM model not only reproduces the dominant flow oscillations but also achieves much smaller and more spatially uniform errors. This confirms that training on structured grids offers a more efficient and accurate surrogate for unsteady wave propagation in ducts.

We observe that the structured mesh model achieves improved accuracy compared with the unstructured mesh model, while also offering slightly reduced training times (24.6 hours as compared to 26.5 hours respectively). This tradeoff suggests that structured discretizations may be advantageous for operator-learning approaches, particularly when scaling to more complex turbomachinery blade-row flows. However, further investigation is required to confirm whether similar gains persist in regions of rapid spatial variation, such as blade wakes or shock–boundary layer interactions, where fine resolution is typically necessary in traditional CFD.

A comparison of Figures 1 and 2 highlights the tradeoffs between the two training strategies. The unstructured mesh model (Figure 1) reproduces the main wave features but exhibits localized discrepancies near the duct outlet (left), where wave steepening and reflections occur. By contrast, the structured mesh model (Figure 2) shows reduced error magnitudes across the duct, particularly in the downstream region, suggesting that the regular discretization aids the network in learning coherent spatial modes. However, both models capture the dominant oscillatory dynamics with comparable fidelity, indicating that the structured mesh provides a more computationally efficient representation without sacrificing predictive accuracy. This efficiency is especially promising for scaling to turbomachinery blade-row problems, where the cost of training on full CFD meshes is prohibitive.

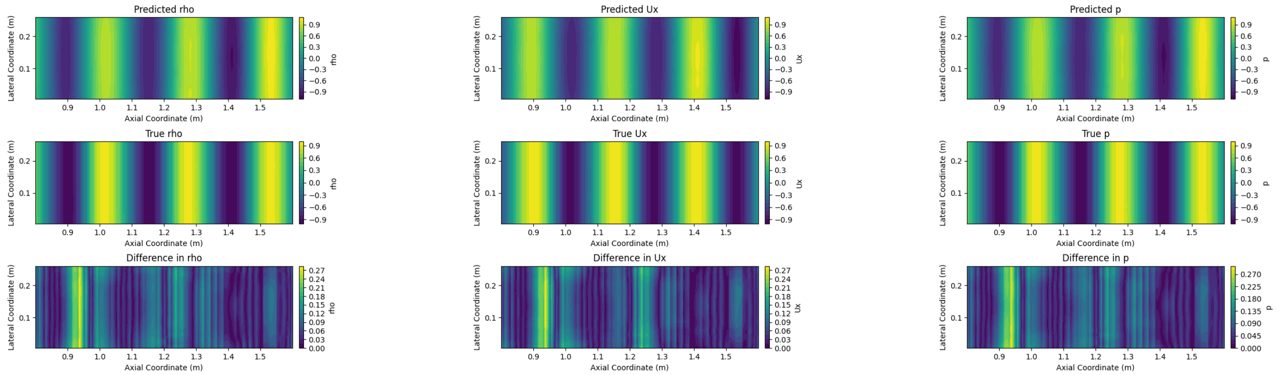

Fig. 2 Predicted flow fields for the final timestep of the last oscillation period on the **structured mesh**(200×20 grid). Results are shown for (a) density ρ, (b) axial velocity U_x, and (c) pressure p. The top row shows the DeepONet–LSTM prediction, the middle row shows the CFD reference solution, and the bottom row shows the absolute difference field.

Finally, we computed normalized error measures across the duct domain for each flow variable, across the final oscillation period, and model variant:

1. **Normalized Root-Mean-Square Error (NRMSE):**

$$\mathrm{NRMSE}_c = \frac{\sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(u_c(x_i) - \hat{u}_c(x_i)\right)^2}}{u_{c,\max} - u_{c,\min}}$$

where u_c and $\hat{u}_c$ denote the CFD and DeepONet–LSTM predictions for flow variable c, and the denominator normalizes by the dynamic range of the variable.

2. **Relative L_2 Error:**

$$E_{L_2,c} = \frac{\|u_c - \hat{u}_c\|_2}{\|u_c\|_2}$$

Table 1 quantifies the predictive accuracy of the DeepONet–LSTM models trained on the unstructured CFD mesh and the structured mesh. On average, the structured mesh reduces the NRMSE from **9.52% to 4.46%** and the relative L_2 error from **0.277 to 0.129**. This represents more than a twofold improvement in accuracy across all flow variables. The gains are consistent for density, velocity, and pressure, indicating that the structured discretization provides a more uniform representation of the spatial dynamics, which the network can exploit to better capture the wave propagation. These results confirm that structured meshes not only accelerate training but also enhance predictive fidelity, making them particularly attractive for extending this framework to more complex turbomachinery configurations.

Table 1 Comparison of error metrics between the DeepONet–LSTM models trained on the unstructured CFD mesh and the structured mesh (200×20 grid).

Variable	**CFD mesh**		**Structured mesh**	
	NRMSE [%]	E_{L_2}	NRMSE [%]	E_{L_2}
Density (ρ)	9.37	0.273	4.49	0.130
Axial velocity (U_x)	9.61	0.279	4.35	0.126
Pressure (p)	9.58	0.279	4.54	0.132
Average	9.52	0.277	4.46	0.129

Conclusion

We have demonstrated the application of a DeepONet architecture with an LSTM-based branch network for predicting unsteady flow fields in an empty duct subjected to sinusoidal outlet forcing. Trained on a single oscillation period, the model accurately reconstructs the flow field during the final period, effectively bypassing thousands of CFD timesteps.

Quantitative error analysis confirms the predictive capability of the approach. When trained on the unstructured CFD mesh, the DeepONet–LSTM model achieved average errors of 9.52% (NRMSE) and 0.277 (L_2). Training on a structured mesh reduced these values by more than half, to 4.46% and 0.129, respectively. The structured discretization thus not only accelerated training but also enhanced predictive fidelity, yielding smoother and more accurate reconstructions of the wave dynamics.

The DeepONet–LSTM framework therefore provides rapid, accurate predictions of unsteady pressure distributions and aerodynamic loads, enabling orders-of-magnitude acceleration relative to conventional CFD. These surrogates can be integrated directly into aeroelastic solvers for flutter margin prediction or leveraged in parametric studies requiring large numbers of operating conditions.

This empty-duct case study serves as a controlled validation of the operator's stability and generalization. Future work will extend the framework to turbomachinery blade-row configurations, where richer unsteady features such as wakes, shocks, and blade interactions must be captured. A natural progression is to benchmark the DeepONet–LSTM model against Fourier Neural Operators (FNOs) [1] on canonical cascade test cases, such as Standard Configuration 10, which supports both subsonic and transonic flow regimes.

Acknowledgments The authors would like to thank Rolls-Royce Plc. for providing the financial support for this project (iCASE #230125) and for giving permission to publish this work. The authors also acknowledge the financial support of the EPSRC (grant #EP/Y528560/1).

References

1. Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., and Anandkumar, A. "Fourier neural operator for parametric partial differential equations" (2021)ifnextchar.gobble.
2. Lu, L., Jin, P., Pang, G., Zhang, Z., and Karniadakis, G.E. "Learning nonlinear operators via deeponet based on the universal approximation theorem of operators". *Nature Machine Intelligence*, 3(3):218–229 (2021)ifnextchar.gobble.
3. He, J., Kushwaha, S., Park, J., Koric, S., Abueidda, D., and Jasiuk, I. "Sequential deep operator networks (s-deeponet) for predicting full-field solutions under time-dependent loads". *Engineering Applications of Artificial Intelligence*, 127:107258 (2024)ifnextchar.gobble.
4. He, J., Kushwaha, S., Park, J., Koric, S., Abueidda, D., and Jasiuk, I. "Predictions of transient vector solution fields with sequential deep operator network". *Acta Mechanica*, 235(8):5257–5272 (2024)ifnextchar.gobble.
5. Loshchilov, I. and Hutter, F. "Decoupled weight decay regularization" (2019)ifnextchar.gobble.