Multi-fidelity Data Reconstruction for Wind Pressure on Building

Yanyu Ke
Department of Civil and Environmental Engineering
The Hong Kong University of Science and Technology
ykeag@connect.ust.hk

Yifan Wang
AIWE
Harbin Institute of Technology, Shenzhen
wangyf@hit.edu.cn

Junle Liu
FLOW, Department of Engineering Mechanics
KTH Royal Institute of Technology
junle@kth.se

Wenliang Chen
AIWE
Harbin Institute of Technology , Shenzhen
wenliang@stu.hit.edu.cn

Gang Hu
AIWE
Harbin Institute of Technology , Shenzhen
hugang@hit.edu.cn

Tim K.T. Tse
Department of Civil and Environmental Engineering
The Hong Kong University of Science and Technology
timkttse@ust.hk

Abstract

Accurate measurement of wind-induced pressures on buildings is crucial yet costly, especially under interference effects where high-fidelity wind tunnel experiments (EXP) remain limited in spatial resolution and sample size. In contrast, computational fluid dynamics (CFD) simulations offer scalable, low-cost data but suffer from systematic biases. Bridging this fidelity gap poses a significant challenge in wind engineering. This study proposes an AI-enhanced framework for experimental data reconstruction, leveraging multi-fidelity data fusion. The focus lies on complex aerodynamic interference scenarios between two square buildings. A comprehensive CFD–EXP dataset spanning 888 configurations is constructed, covering multiple building layouts and incident wind angles. The sparse-to-dense reconstruction is performed, where limited high-fidelity data are used to enhance predictions. Sparse experimental sensors are incorporated into multi-fidelity neural network (MFNN) and neural operator (MFNO) frameworks: MFNN provides case-specific reconstructions, while MFNO generalizes across building spacings and wind angles without retraining per case. Experimental results confirm significant improvements in spatial reconstruction under sparse supervision. The proposed framework demonstrates strong generalization and robustness. This work advances AI-enhanced experimental data reconstruction, reducing testing costs while enhancing prediction reliability in wind engineering.

39th Conference on Neural Information Processing Systems (NeurIPS 2025).

1 Introduction

The accurate assessment of wind loads on building surfaces is a cornerstone of structural design, ensuring the safety, serviceability, and resilience of civil infrastructure against extreme aerodynamic forces [Blocken, 2014]. In wind engineering, computational fluid dynamics (CFD) simulations and wind tunnel experiments (EXP) are the two primary methods for obtaining wind pressure data on building surfaces [Kareem, 2020, Kwok and Hu, 2023]. CFD is widely used due to its flexibility and low cost, and it can generate high-resolution pressure fields under various configurations [Ding and Kareem, 2018, Chen et al., 2024]. However, the accuracy of CFD results can be significantly affected by modeling assumptions [Zhang et al., 2020]. In contrast, wind tunnel experiments provide more reliable and physically accurate data [Li and Li, 2017], but the measurements are limited to discrete locations and require significant time and cost to obtain [Ke et al., 2022, Liu et al., 2024b].

Traditional methods such as linear or polynomial interpolation often fail to capture the complex, nonlinear behavior of wind-induced pressure fields, especially under flow separation and aerodynamic interference [Hu et al., 2020, Zhou et al., 2021, Chen et al., 2025]. The emerging machine learning models, such as multilayer perceptrons [Tian et al., 2020], convolutional neural networks [Guo et al., 2016], or other advanced techniques [Bre et al., 2018, Hu and Kwok, 2020, Gao et al., 2024, Liu et al., 2024a] have shown promising results in aerodynamic prediction, but their reliance on single-fidelity data limits their robustness and generalization to unseen configurations or sparse inputs.

To address these challenges, two major research directions have emerged in recent years. Neural operators have been proposed to learn mappings between function spaces rather than finite-dimensional vectors [Kovachki et al., 2023, Li et al., 2021]. Notably, Lu et al. [2021] proposed the DeepONet that enables strong generalization across spatial domains and parametric variations by decoupling functional inputs from specific discretizations. In parallel, multi-fidelity neural network (MFNN) has been developed to integrate low-fidelity simulations and high-fidelity measurements. Meng and Karniadakis [2020] introduced hybrid architectures combining linear and nonlinear mappings to fuse low-fidelity and high-fidelity sources. Furthermore, Wang et al. [2025] utilized a novel multi-fidelity neural network to enhance the model's applicability to nonlinear inconsistency issues, which are commonly seen in transonic aerodynamic problems. In an effort to improve generalization, Howard et al. [2023] further extended neural operator architectures into multi-fidelity settings, enabling function-level fusion across spatial resolutions. These advances provide promising tools for wind pressure reconstruction, enabling models that are both generalizable and data-efficient.

In this work, we propose a dual framework that integrates CFD simulations and wind tunnel experiments to predict surface pressure on buildings. When sparse experimental data is available, we fuse it with CFD simulations using multi-fidelity learning. We adopt multi-fidelity neural networks (MFNN) for case-specific training, and multi-fidelity neural operators (MFNO) for generalizable learning across cases without retraining. We validate our approach on a dataset of 888 CFD–EXP configurations involving aerodynamic interference between twin square buildings. Dense experimental measurements serve as the reference, while sparse subsets are used to test the framework's ability to reconstruct full-field surface pressures.

2 Data Generation and Collection

This study investigates aerodynamic interference in two buildings. As shown in Figure 1a, the main building is at the origin. The interference building offers 37 discrete positions, and we also vary the incident wind direction for aerodynamic characterization. Accordingly, we assembled a dataset of 888 cases spanning 37 configurations and 24 wind directions ($0°$–$345°$ at $15°$ increments) using an automated modeling framework. Position indices are consistent with the Tokyo Polytechnic University (TPU) aerodynamic database convention [1].In this study, two primary data sources were employed to obtain wind pressure distributions on building surfaces, particularly under interference effects between two buildings. The first method involves computational fluid dynamics (CFD) simulations, while the second method uses experimental wind tunnel data (EXP).

For CFD data, all simulations were performed with a uniform inlet velocity of 1 m/s. The square building section in the simulation had a characteristic length B of 28 meters, which defines the

[1]For more information and details of the TPU database, please visit `https://wind.arch.t-kougei.ac.jp/system/eng/contents/code/tpu`

full scale of the building. All computations were carried out using the open-source CFD library *OpenFOAM*. A two-dimensional (2D) configuration was adopted to reduce computational cost relative to 3D while retaining the primary interference mechanisms [Liu et al., 2024a]. Further details of the automated model-generation pipeline and the numerical setups are provided in Appendix A.

For EXP data, the wind tunnel model was scaled at a 1/400 length scale. The principal building, which had pressure taps distributed across its surface (70 mm breadth, 70 mm depth, and 280 mm height in 3D model scale). A total of 252 pressure taps were distributed on the building surfaces, covering all four sides of the building. It was divided into 9 horizontal layers, representing the 9 different heights of the building. Each layer contained 28 pressure taps, with 7 taps placed along each row on every side of the building, as shown in Figure 1b

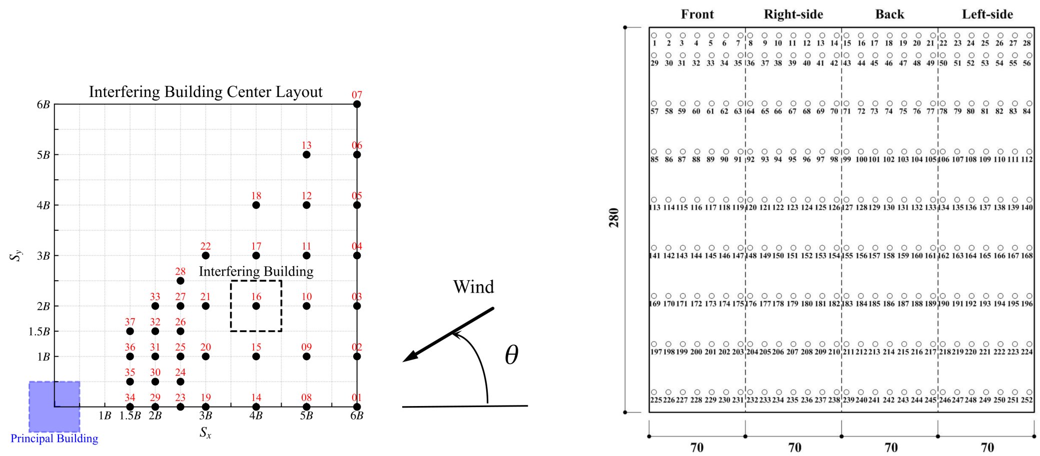

(a) Locations of interference buildings.

(b) Pressure points definition.

Figure 1: Locations of interference buildings, wind angle and pressure points definition.

3 Results and Discussions

3.1 Multi-Fidelity Sparse-to-Dense Reconstruction

To reduce reliance on large experimental datasets and enable robust reconstruction under sparse supervision, the Multi-Fidelity Neural Network (MFNN) framework is introduced, further details of the model framework are provided in Appendix B. The objective is to reconstruct the full high-fidelity experimental distribution using only a small number of experimental measurements together with low-fidelity CFD data. In practice, full-field experimental measurements are rarely available due to sensor limitations. To address this, we combine dense CFD simulations with sparse experimental data for reconstruction via MFNN.

Low-fidelity neural network (NN_L) is first trained on 28 CFD points to provide a low-fidelity pressure field. Its predictions at 8 sparse sensor locations, together with their coordinates, are then corrected using the corresponding experimental data, yielding an accurate reconstruction of the full-field pressure distribution. As shown in Figure 2a, the reconstructed results match well with the experimental distribution, including regions without direct observations. The unused 20 experimental points serve as independent validation, confirming that the MFNN generalizes beyond the observed subset. This setup mimics realistic scenarios where only limited measurements are available and ensures that reconstruction performance reflects genuine sparse-to-full generalization. MFNN provides a practical way to exploit low-fidelity CFD fields while calibrating them with minimal experimental input, enabling accurate full-field predictions under highly constrained measurement conditions.This reconstruction not only matches the experiment well, but also clearly outperforms cubic interpolation, which fails to capture the abrupt pressure reversals caused by flow separation and reattachment. This highlights the effectiveness of AI-assisted sparse measurement recovery in reducing instrumentation requirements while maintaining accuracy.

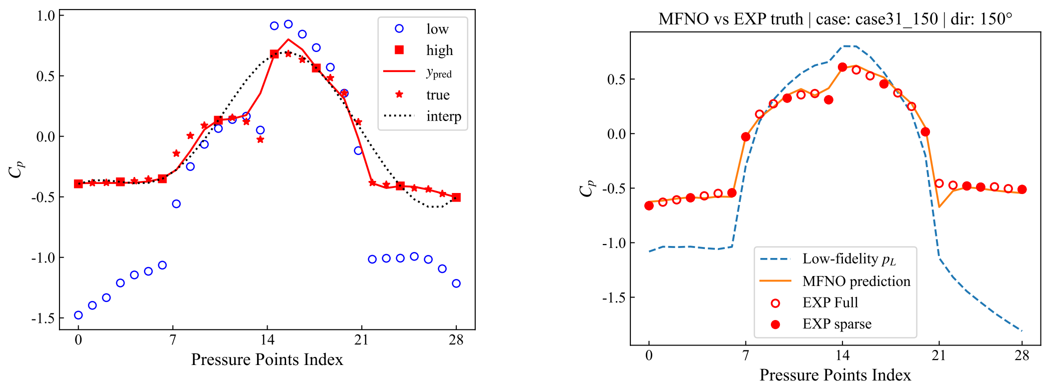

(a) MFNN reconstructed prediction shows close agreement with experimental measurements (high and true points), validating accurate full-field recovery from sparse sensors.

(b) MFNO reconstruction for a sample case, showing close agreement with sparse training sensors and independent validation data.

Figure 2: MFNN and MFNO reconstruction

3.2 Generalized Multi-Fidelity Neural Operator

To enhance generalization, the MFNN framework is extended to all 888 cases spanning diverse upstream layouts and wind directions. The low-fidelity neural operator NN_L is trained on 888×28 CFD points to capture shared flow features across the dataset, while the correction network leverages 888×12 experimental observations with their low-fidelity counterparts. Further details of the model framework (Multi-Fidelity Neural Operator) are provided in Appendix C.

To better illustrate the progression from MFNN to MFNO, Table 1 provides a side-by-side comparison, emphasizing the enhanced scope, training strategy, and generalization capacity of MFNO.

Table 1: Comparison between MFNN and MFNO.

Aspect	**MFNN**	**MFNO**
Scope	Single-case	Multi-case
Inputs	CFD pressures + sparse EXP	CFD pressures and sparse EXP per case
Supervision	Few EXP points	Few EXP points per case
Generali-zation	Limited to the given case	Transferable across different cases
Output	Full-field EXP reconstruction for one case	Scalable calibrated fields for unseen cases

As depicted in Figure 2b, an example case (case31 at $150°$) is shown to illustrate the reconstruction performance. A small subset of experimental points is designated as training labels for the correction network, while the remaining is held out as independent validation. The MFNO predictions align well with the EXP sparse and EXP full, substantially correcting the bias in the raw low-fidelity neural operator. This demonstrates that the framework achieves strong generalization across diverse configurations, delivering accurate reconstructions without retraining, and thus providing a scalable surrogate for experimental measurements. More cases of MFNO performance are demonstrated in Appendix C.

4 Conclusion

This study introduces an AI-enhanced framework for wind pressure prediction and experimental data reconstruction in complex urban aerodynamic scenarios. By combining CFD simulations with wind tunnel experiments, the sparse-to-dense multi-fidelity learning is developed, which integrates sparse experimental measurements with dense CFD simulations to recover full-field high-fidelity pressure distributions, remaining accurate even under limited sensor availability. Validation on 888 CFD–EXP configurations demonstrates both robustness and strong generalization, enabling accurate predictions without retraining across diverse layouts and wind directions. By enabling low-cost, high-accuracy estimation of wind-induced pressures under interference conditions, this work contributes to the advancement of data-efficient and generalizable aerodynamic modeling.

The proposed framework provides a clear path toward deployment in wind engineering practice. In wind tunnel testing, the multi-fidelity sparse-to-dense reconstruction reduces the number of sensors required while still recovering full-field pressures with high accuracy. In real-world projects, where sensor installations are even more limited, the framework can leverage CFD-informed corrections to provide reliable estimates. The demonstrated generalization across 888 configurations indicates strong potential for urban-scale aerodynamic assessments, supporting structural safety evaluation, wind-resistant design, and digital twin development for the built environment.

References

B. Blocken. 50 years of computational wind engineering: Past, present and future. *Journal of Wind Engineering and Industrial Aerodynamics*, 129:69–102, 2014.

F. Bre, J. M. Gimenez, and V. D. Fachinotti. Prediction of wind pressure coefficients on building surfaces using artificial neural networks. *Energy and Buildings*, 158:1429–1441, 2018.

L. Chen, C. Li, J. Wang, G. Hu, and Y. Xiao. A coherence-improved and mass-balanced inflow turbulence generation method for large eddy simulation. *Journal of Computational Physics*, 498: 112706, 2024.

W. Chen, Z. Wang, H. Hong, J. Song, and G. Hu. Aerodynamic interference effects on three connected high-rise buildings with y-plan layout. *Engineering Structures*, 326:119494, 2025.

F. Ding and A. Kareem. A multi-fidelity shape optimization via surrogate modeling for civil structures. *Journal of Wind Engineering and Industrial Aerodynamics*, 178:49–56, 2018.

H. Gao, G. Hu, D. Zhang, W. Jiang, K. T. Tse, K. Kwok, and A. Kareem. Urban wind field prediction based on sparse sensors and physics-informed graph-assisted auto-encoder. *Computer-Aided Civil and Infrastructure Engineering*, 39(10):1409–1430, 2024.

X. Guo, W. Li, and F. Iorio. Convolutional neural networks for steady flow approximation. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 481–490, San Francisco, CA, USA, 2016. ACM.

A. A. Howard, M. Perego, G. E. Karniadakis, and P. Stinis. Multifidelity deep operator networks for data-driven and physics-informed problems. *Journal of Computational Physics*, 493:112462, 2023.

G. Hu and K. Kwok. Predicting wind pressures around circular cylinders using machine learning techniques. *Journal of Wind Engineering and Industrial Aerodynamics*, 198:104099, 2020.

G. Hu, L. Liu, D. Tao, J. Song, K. T. Tse, and K. Kwok. Deep learning-based investigation of wind pressures on tall building under interference effects. *Journal of Wind Engineering and Industrial Aerodynamics*, 201:104138, 2020.

A. Kareem. Emerging frontiers in wind engineering: Computing, stochastics, machine learning and beyond. *Journal of Wind Engineering and Industrial Aerodynamics*, 206:104320, 2020.

Y. Ke, G. Shen, X. Yang, and J. Xie. Effects of surface-attached vertical ribs on wind loads and wind-induced responses of high-rise buildings. *Sustainability*, 14(18):11394, 2022.

N. B. Kovachki, Z. Li, B. Liu, K. Azizzadenesheli, K. Bhattacharya, A. M. Stuart, and A. Anandkumar. Neural operator: Learning maps between function spaces with applications to partial differential equations. *Journal of Machine Learning Research*, 24(92):1–96, 2023.

K. Kwok and G. Hu. Wind energy system for buildings in an urban environment. *Journal of Wind Engineering and Industrial Aerodynamics*, 234:105349, 2023.

S. Li and M. Li. Spectral analysis and coherence of aerodynamic lift on rectangular cylinders in turbulent flow. *Journal of Fluid Mechanics*, 830:408–438, 2017.

Z. Li, N. B. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. M. Stuart, and A. Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv:2010.08895*, 2021.

J. Liu, K. Shum, K. T. Tse, and G. Hu. Bidirectional prediction between wake velocity and surface pressure using deep learning techniques. *Physics of Fluids*, 36(2):025162, 02 2024a.

J. Liu, K. T. Tse, G. Hu, C. Liu, B. Zhang, and K. C. S. Kwok. Exploring aerodynamics of a rectangular cylinder using flow field and surface pressure synchronized testing technique. *Physics of Fluids*, 36(8):085174, 08 2024b.

L. Lu, P. Jin, and G. E. Karniadakis. Deeponet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3:218–229, 2021.

X. Meng and G. E. Karniadakis. A composite neural network that learns from multi-fidelity data: Application to function approximation and inverse pde problems. *Journal of Computational Physics*, 401:109020, 2020.

J. Tian, K. R. Gurley, M. T. Diaz, P. L. Fernández-Cabán, F. J. Masters, and R. Fang. Low-rise gable roof buildings pressure prediction using deep neural networks. *Journal of Wind Engineering and Industrial Aerodynamics*, 196:104026, 2020.

X. Wang, H. Li, H. Lin, H. Tang, and W. Zhang. Multisource aerodynamic data reconstruction method using an enhanced multifidelity neural network. *Engineering Applications of Artificial Intelligence*, 159:111707, 2025.

X. Zhang, A. U. Weerasuriya, and K. T. Tse. Cfd simulation of natural ventilation of a generic building in various incident wind directions: Comparison of turbulence modelling, evaluation methods, and ventilation mechanisms. *Energy and Buildings*, 229:110516, 2020.

L. Zhou, K. T. Tse, G. Hu, and Y. Li. Mode interpretation of interference effects between tall buildings in tandem and side-by-side arrangement with pod and ica. *Engineering Structures*, 243: 112616, 2021.

A Simulation Process Details

The CFD data was simulated using the open-source computational fluid dynamics library, *OpenFOAM*. The model considered cases in two dimensions (2D), which contributed to its relatively fast processing speed compared to 3D models. Unsteady Reynolds-averaged Navier-Stokes (URANS) simulations were employed to model the flow field around two buildings and to calculate the wind pressure on the building surfaces. The *pimplefoam* algorithm was used to solve the velocity-pressure coupling problems. Temporal discretization was carried out using a second-order implicit scheme, while spatial discretization was performed using the *Gaussian linear* upwind scheme. The boundary conditions were as follows: the inlet boundary was set as a velocity inlet with a flow speed of 1 m/s, while the outlet boundary was defined as a pressure outlet. Symmetry boundary conditions were applied to the lateral boundaries, and a non-slip fixed wall condition was used for the building surfaces. To facilitate reproducibility and clarity, the automated model-generation pipeline is described in detail in the following subsection.

I. Model Generation

The dataset for deep learning model training is generated through an automated simulation process. *STL* models of the two bluff bodies configuration are automatically generated using signed distance functions. The simulation process in *OpenFOAM* is fully automated by executing shell scripts, which control all aspects of the workflow. The dataset for deep learning model training is generated through an automated simulation process. *STL* models of the two bluff bodies configuration are automatically generated using signed distance functions. The simulation process in *OpenFOAM* is fully automated by executing shell scripts, which control all aspects of the workflow.

II. Geometric Information Extraction

The geometric details of the models are extracted using the *SurfaceFeatures* tool within *OpenFOAM*. This step ensures the correct representation of the model geometry, which is critical for accurate simulation results.

III. OpenFOAM Setup

The primary simulation parameters are set in *OpenFOAM* files. These include:

- Solver: *pimplefoam* algorithm
- Turbulence model: SST k-ω model
- Time discretization: second-order implicit scheme
- Spatial discretization: *Gauss linear* upwind scheme
- Mesh parameters: time step, mesh division, etc.

These parameters are kept constant throughout the entire simulation loop to ensure consistency and repeatability. The *OpenFOAM* configuration is illustrated in Table A1.

IV. CFD Simulation and Post-Processing

Once the CFD control equations are solved, post-processing is performed to calculate the necessary fluid dynamics parameters. The resulting flow field data is then stored as datasets for the training of the deep learning model, enabling rapid prediction of flow fields around the bluff bodies.

V. Mesh Generation Using SnappyHexMesh

During the simulation process, the *SnappyHexMesh* tool is used for mesh generation. This tool, built into *OpenFOAM*, creates meshes that fit the model geometry through an iterative process. The steps involved in mesh generation are:

- Creating the Background Mesh: The *blockMesh* tool is used to generate a hexahedral background mesh.
- Extracting Geometric Features: The *surfaceFeatures* tool is applied to extract geometric features from the model.
- Setting *SnappyHexMeshDict* Parameters: Basic parameters are configured in the *SnappyHexMeshDict* file. These settings are critical for determining the mesh quality. Mesh refinement is applied in specific regions, particularly around the building surfaces, to enhance computational accuracy.
- Running *SnappyHexMesh* in Parallel: The *SnappyHexMesh* process is executed in parallel, ensuring efficient mesh generation.

B Multi-Fidelity Neural Network (MFNN) for Sparse Reconstruction

To reduce reliance on large experimental datasets and enable robust reconstruction under sparse supervision, the Multi-Fidelity Neural Network (MFNN) framework is introduced(shown in Figure A1). The objective is to reconstruct the full high-fidelity experimental distribution using only a small number of experimental measurements together with low-fidelity CFD data.

Given CFD predictions at m surface points

$$\boldsymbol{p}_L = [p(x_{L,1}), p(x_{L,2}), \dots, p(x_{L,m})] \in \mathbb{R}^m,$$

we first build a low-fidelity surrogate $\mathcal{NN}_L$ that approximates the CFD pressure field. A sparse set of k experimental samples

$$\{(x_{H,i}, p_H(x_{H,i}))\}_{i=1}^{k}, \quad k \ll m,$$

is then used for supervision. At these sparse locations, the model combines $\boldsymbol{p}_L(x_{H,i})$ with coordinates $x_{H,i}$ and applies correction networks to predict experimental pressure. Specifically, MFNN employs

Table A1: OpenFOAM simulation setup.

Parameter	**Value**
Cross-section size (B)	28 m
Inlet boundary	Velocity inlet
Reference (inlet) speed	1 m/s
Outlet boundary	Pressure outlet
Lateral boundaries	Symmetry
Walls	Nonslip wall
Turbulence model	URANS (SST k-ω)
Numerical solver	*pimpleFoam*
V–p coupling	*PIMPLE*-based
Time discretization	Second-order implicit
Spatial discretization	*Gauss linear* Upwind scheme

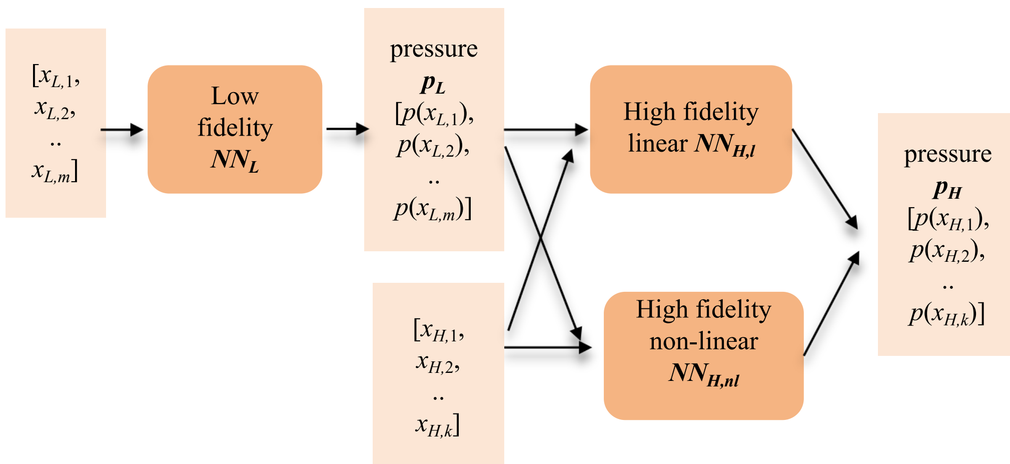

Figure A1: MFNN for sparse-to-full reconstruction. Low-fidelity CFD predictions p_L are combined with sparse experimental measurements to train linear and nonlinear correction networks, producing the full experimental distribution $\boldsymbol{p}_H$.

both a linear mapping $NN_{H,l}$and a nonlinear mapping $NN_{H,nl}$, yielding the reconstruction as follows:

$$\boldsymbol{p}_H(x) = NN_{H,l}(x, \boldsymbol{p}_L(x)) + NN_{H,nl}(x, \boldsymbol{p}_L(x)),$$

where $NN_{H,l}$ captures the dominant linear correlation between CFD and EXP, such as global scaling or bias, while $NN_{H,nl}$ accounts for residual nonlinear discrepancies induced by turbulence modeling or local boundary effects. By combining these two components, MFNN achieves both stability and flexibility in reconstructing the high-fidelity field from sparse supervision. Training is performed only on the k sparse points, while the remaining $(m - k)$ experimental data are reserved for validation. MFNN exploits low-fidelity CFD fields while calibrating them with minimal experimental input, enabling accurate full-field predictions under highly constrained measurement conditions.

C MFNO: Multi-case Generalization over various Configurations

To further demonstrate generalization, we extend the multi-fidelity framework to a neural operator (MFNO), as illustrated in Figure A2. In contrast to the MFNN model, which targets a single configuration, MFNO is trained on a large dataset of 888 cases covering diverse interference locations and wind directions. This allows the model to learn transferable corrections across configurations rather than being restricted to one case.

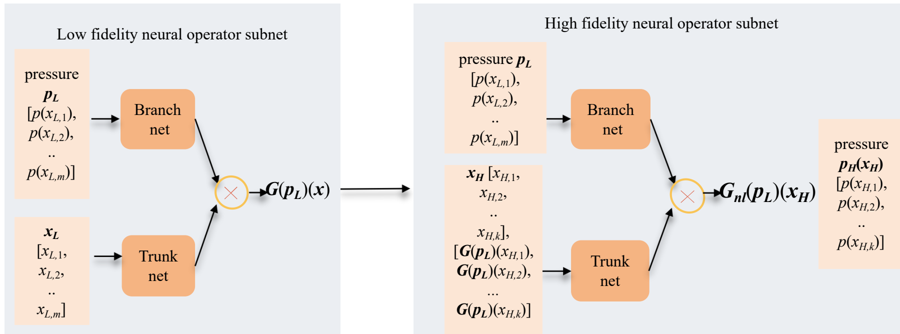

Figure A2: MFNO framework for multi-case generalization. The model learns a neural operator that calibrates low-fidelity CFD pressures $\boldsymbol{p}_L$ to experimental distributions $\boldsymbol{p}_H$ across N diverse configurations.

The input consists of N (case sample number) $\times$ m (low-fidelity points) CFD pressure values

$$\boldsymbol{p}_L = [p(x_{L,1}), p(x_{L,2}), \ldots, p(x_{L,m})],$$

which are first mapped by a low-fidelity operator subnet $\boldsymbol{G}(\boldsymbol{p}_L)(x)$. For each configuration, a sparse set of k experimental samples $\{(x_{H,i}, p_H(x_{H,i}))\}_{i=1}^{k}$ is provided as supervision. The high-fidelity operator subnet then takes $\boldsymbol{p}_L$, the coordinates $\boldsymbol{x}_H$, and the high-fidelity outputs $\boldsymbol{G}_{nl}(\boldsymbol{p}_L)(x_H)$ at these locations, and produces the corrected prediction

$$\boldsymbol{p}_H(\boldsymbol{x}_H) = \boldsymbol{G}_{nl}(\boldsymbol{p}_L)(\boldsymbol{x}_H), \quad \boldsymbol{x}_H \in \{x_{H,1}, \ldots, x_{H,k}\}.$$

Training is carried out only on the k supervised points per case, while the remaining $(m-k)$ experimental points are reserved for validation. Unlike MFNN, which focuses on a single configuration, MFNO aggregates information from $N = 888$ cases covering diverse upstream layouts and wind directions. More cases of MFNO performance are demonstrated as follows in Figure A3:

This large-scale multi-fidelity training enables the operator to learn both global linear correlations and residual nonlinear discrepancies between CFD and EXP in a transferable manner. As a result, MFNO achieves robust and data-efficient calibration of CFD simulations, generalizing across previously unseen geometries and flow conditions.

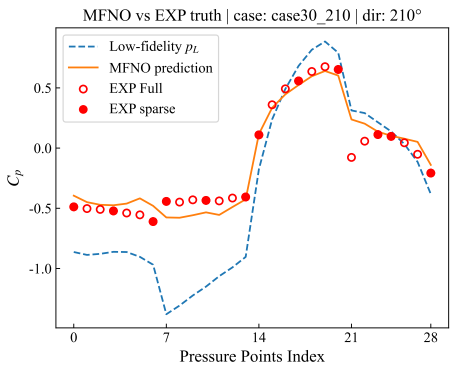

(a) case30 at $(3B, 0.5B)$ at 210°

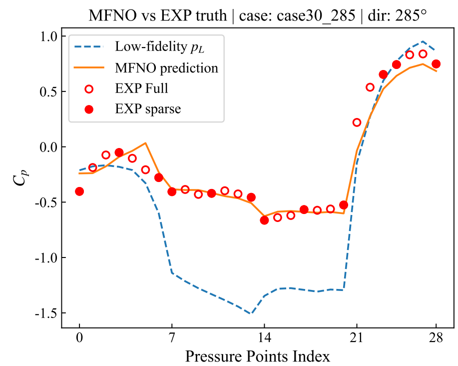

(b) case30 at $(3B, 0.5B)$ at 285°

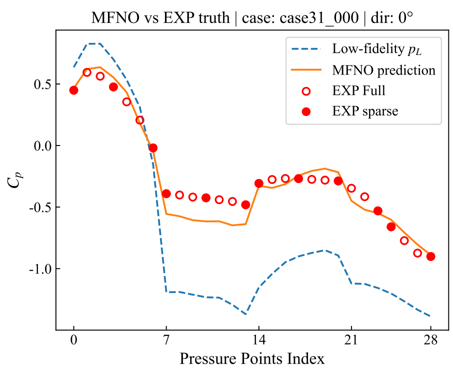

(c) case30 at $(3B, B)$ at 0°

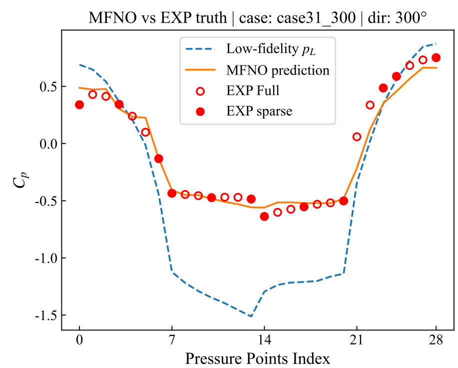

(d) case30 at $(3B, B)$ at 300°

Figure A3: MFNO reconstruction performance

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification:

   Guidelines:

   - The answer NA means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification:

   Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [NA]

   Justification:

   Guidelines:

   - The answer NA means that the paper does not include theoretical results.
   - All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
   - All assumptions should be clearly stated or referenced in the statement of any theorems.
   - The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
   - Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
   - Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification:

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.

- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [No]

   Justification: It will be open access upon publication.

   Guidelines:

   - The answer NA means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: No available now.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification:

   Guidelines:

   - The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification:

    Guidelines:

    - The answer NA means that there is no societal impact of the work performed.
    - If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
    - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
    - The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
    - If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

    Answer: [NA]

    Justification:

    Guidelines:

    - The answer NA means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.

- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [NA]

    Justification:

    Guidelines:

    - The answer NA means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
    - If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [NA]

    Justification:

    Guidelines:

    - The answer NA means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [NA]

    Justification:

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [NA]

    Justification:

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

    Answer: [NA]

    Justification:

    Guidelines:

    - The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy (`https://neurips.cc/Conferences/2025/LLM`) for what should or should not be described.