PresCast: Physics-Constrained Fourier Kolmogorov-Arnold Networks for Bluff Bodies Spatiotemporal Wall Pressure Forecast

Junle Liu
Department of Civil & Environmental Engineering
Hong Kong University of Science and Technology
jliueb@connect.ust.hk

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

Haoyan Li
Aero-Thermo-Mechanics Laboratory
Universite Libre de Bruxelles
l.haoyan.gm@gmail.com

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

Tianle Niu
School of Future Technology
Harbin Institute of Technology
2023110605@stu.hit.edu.cn

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

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

Abstract

Understanding and forecasting wall pressure on bluff bodies is crucial for assessing urban structural safety and infrastructure resilience. While recent studies have explored deep learning for predicting mean and fluctuating pressure coefficients, most existing frameworks are restricted to single-snapshot predictions and lack physical interpretability. This work introduces PresCast, a physics-constrained Fourier Kolmogorov-Arnold neural network (FKAN) designed for spatiotemporal wall pressure forecasting. FKAN integrates a Fourier neural operator within the Kolmogorov-Arnold architecture, and incorporates a physics-constrained frequency loss to guide the training process and enhance generalization to high-frequency dynamics. Wind tunnel experiments were conducted on a classical bluff body, rectangular cylinder of side ratio 1.5 under zero angle of attack, using high-frequency pressure scanning, building a comprehensive dataset for training and validation. Results demonstrate that PresCast achieves accurate spatiotemporal forecasting of wall pressure, capturing both statistical properties and physical signatures, including temporal fluctuations, power spectral density, and spatial distribution. This study highlights the potential of physics-constrained deep learning frameworks for advancing urban aerodynamics and resilient infrastructure design.

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

1 Introduction

Bluff body wall pressure plays an essential role in understanding the urban structure safety, urban infrastructure resilience, and urban building physics. Typically, wall pressure information collection is based on either wind tunnel testing or numerical simulations, which are time-consuming and costly. Recent advances in employing artificial intelligence (AI) algorithms to forecast or generate wall pressure information in some unseen conditions have received great attention [Hu et al., 2020a, Weng and Paal, 2022], as historical research has accumulated a large database providing a solid foundation for artificial intelligence model training and validation [Ho et al., 2005, Pierre et al., 2005, Oh et al., 2007, Li et al., 2024, Liu et al., 2025]. For example, Fernández-Cabán et al. [2018] built an artificial neural network (ANN) integrating backpropagation training to predict the mean, root mean square, and peak pressure coefficients on three scaled rectangular cylinders. Experimental data were collected to train the model. Their evaluations presented that the ANN model was capable of predicting accurate mean pressure coefficients compared to the experimental data. Tian et al. [2020] developed a deep neural network to predict mean and peak wall pressure coefficients on rectangular cylinder-shaped buildings. Based on the experimental data collected, the trained deep neural network can make satisfactory predictions on the mean pressure coefficients.

However, existing AI-based pressure predictions have focused mainly on using full-dimensional data to predict the instantaneous wall pressure distribution [Liu et al., 2024a], the mean pressure distribution [Weng and Paal, 2022, Liu et al., 2025], the fluctuation pressure coefficient [Hu et al., 2020b], and the maximum pressure coefficient [Tian et al., 2020]. Few works have focused on spatiotemporal wall pressure forecast. In particular, there is an obvious absence of spatiotemporal wind tunnel wall pressure measurements forecast. This is partly due to the challenges posed by wind tunnel data, such as measurement sparsity, noise contamination, and non-stationary flow characteristics [Chen et al., 2021, Santos et al., 2023, Li et al., 2025], which make spatiotemporal and high-fidelity forecasting more difficult.

To address the above challenges, we propose a bluff body wall pressure forecast framework (PresCast), namely the Fourier Kolmogorov-Arnold neural network (FKAN) deep learning model, to make spatiotemporal wall pressure predictions and generation based on a collected historical dataset. FKAN is developed based on the Fourier neural network architecture [Li et al., 2020, 2023] and coupled with the Kolmogorov-Arnold neural Network theory [Liu et al., 2024c]. In this work, the general layout is as follows: Section 2 introduces the problem definition and FKAN model details, and Section 3 presents the dataset collection and model training settings. Finally, the results are discussed in Section 4.

2 Problem setup and methodology

We consider the problem of forecasting time-series bluff body wall pressure with both temporal and spatial dependencies. The collected bluff body wall pressure information is defined as

$$\mathbf{P}_{in} \in \mathbb{R}^{T \times C_{\text{in}}},$$

where $T = 1024$ is the number of time-series snapshots for the input information, i.e., time instant index $t \in (i, i + 1024)$, and $C_{\text{in}} = 26$ is the number of spatial points on the bluff body walls. The forecast time-series wall pressure information has the same shape but in different time instants index as:

$$\mathbf{P}_f \in \mathbb{R}^{\Delta T \times C_{\text{out}}},$$

where $C_{in} = C_{out}$ denotes the identical spatial domains, ΔT is the forecast time-series interval starting from the last snapshot of $\mathbf{P}_{in}$, i.e., time instant index $t \in (i + 1025, i + 2048)$. In this work, we set $\Delta T = T$, indicating the forecast time length is identical to the input time length as 1024. We develop a Fourier Kolmogorov-Arnold neural Network (FKAN) deep learning model, namely $\boldsymbol{\mathcal{G}}_\theta$, to forecast future time instant wall pressure $\mathbf{P}_f$ based on historical wall pressure $\mathbf{P}_{in}$ as:

$$\mathbf{P}_f = \boldsymbol{\mathcal{G}}_\theta(\mathbf{P}_{in}) \tag{1}$$

The FKAN model consists of several modules, including the Fourier Neural Network [], convolutional operator, and KAN activation. Module details of the FKAN model and mathematical expressions are presented in Appendix A.

3 Data Generation and model training

The experiment was carried out in a closed-loop wind tunnel with a test section of $800\ mm \times 500\ mm \times 500\ mm$ (length $\times$ width $\times$ height). In this experiment, the approaching flow velocity is set at 10 ***m/s***. The turbulence intensity of the wind tunnel is less than 0.4% evaluated with the ***Hanghua CTA04-EDU*** hot-wire, indicating that the approaching flow is treated as uniform flow [Liu et al., 2024b]. High-frequency pressure scanning with sampling frequency f_s=400 Hz is used to sample wall pressure on 26 pressure taps on rectangular cylinder surface. The rectangular cylinder tested is in the side ratio = 1.5 under the angle of attack = 0°. The total number of temporal snapshots collected for the wall pressure is 24,000. The pressure taps arrangement, rectangular cylinder dimension information, experimental devices, and data validation are shown in Appendix B.

After collecting 24,00 instantaneous wall pressure snapshots on 26 pressure taps, the dataset is processed and prepared for model training and testing. The dataset of size (24,000, 26) is first standardized to (0, 1) for better training convergence. In the data sampling part, we sample the data in time sequence according to the target. Data mapping strategy is presented in the Appendix B.

In the model training part, we design the loss control as two components as follows:

$$\mathcal{L} = \mathcal{L}_1 + \beta \mathcal{L}_f, \tag{2}$$

where $\mathcal{L}_1$ is the mean square error (MSE) loss for the experimental ground truth and forecast results, and $\mathcal{L}_f$ is the frequency discrepancies calculated from the forecast results and experimental ground truth, β is a dynamic coefficient to adjust the physical loss defined as follows:

$$\beta = \beta_{end} \frac{epo}{epochs}, \tag{3}$$

where $\beta_{end} = 0.001$ is a fixed value to adjust the model concentration on physical information, epo is the current training epoch, and $epochs$ is total epoch set for the model training. As for the physical frequency discrepancies $\mathcal{L}_f$, we performed a fast Fourier transform $\mathcal{F}$ on the time-series wall pressure $\mathbf{P}_j$ on the j^{th} spatial pressure tap for both experimental ground truth and forecast results as:

$$\mathbf{M}_{j,in} = \mathcal{F}(\mathbf{P}_{j,in}), \tag{4}$$

$$\mathbf{M}_{j,f} = \mathcal{F}(\mathbf{P}_{j,f}), \tag{5}$$

where $\mathbf{M}_j$ denotes the complex vector of the time-series wall pressure at j^{th} pressure tap. Then we calculated the frequency discrepancies based on the Fourier transform vectors norms as follows:

$$\mathcal{L}_f = \mathbb{E}\left[\|\mathbf{M}_{j,in} - \mathbf{M}_{j,f}\|_2^2\right]. \tag{6}$$

More details on model training hyperparameters settings and device usage are shown in Appendix C.

4 Results

We evaluate the results from two main perspectives, including statistical comparison and physical verifications. Based on the forecast pressure information, we evaluate the statistical information between the forecast data and experimental ground truth, such as time-averaged wall pressure information $\overline{C}_p$, standard deviation of wall pressure fluctuation σ on every single pressure tap shown in Figure 1a, standard deviation value of lift coefficient $C_{l,std}$, and time-averaged drag coefficient $\overline{C}_d$ shown in Figure 1b. More details on calculating statistical information can be found in Appendix B. As indicated by the comparison, we can see that the forecast results are in a good alignment with the experimental ground truth.

Apart from the statistical information, we compare the physical information between the forecast results and experimental ground truth from an instantaneous perspective and a time-series perspective. The instantaneous wall pressure distribution of $\Delta T/\delta t = 983$ is compared on all 26 pressure taps shown in Figure 2a and the time history comparison of the entire 1024 snapshots of pressure tap 8 is shown in Figure 2b. In addition, we compare the time history of the instantaneous lift coefficient C_l (Figure 2c) and the instantaneous drag coefficient C_d (Figure 2d) between forecast results and experimental data. From the results shown in Figure 2, we can derive that the FKAN model has the capability to forecast time-series wall pressure on the tested bluff body.

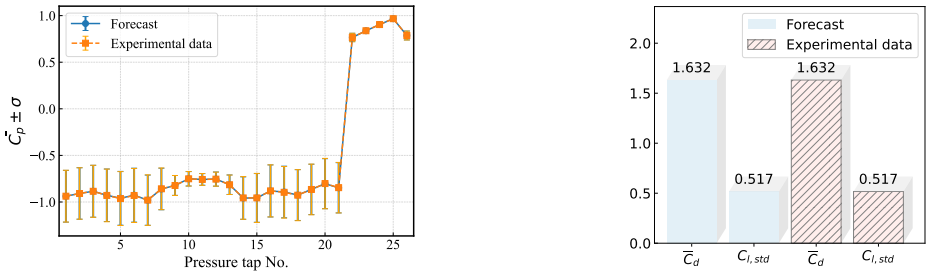

(a) (b)

Figure 1: Statistical information comparison between forecast results and experimental ground truth: (a) Time-averaged wall pressure coefficient $\overline{C}_p$ and standard deviation (σ) on each pressure tap, (b) Time-averaged drag .

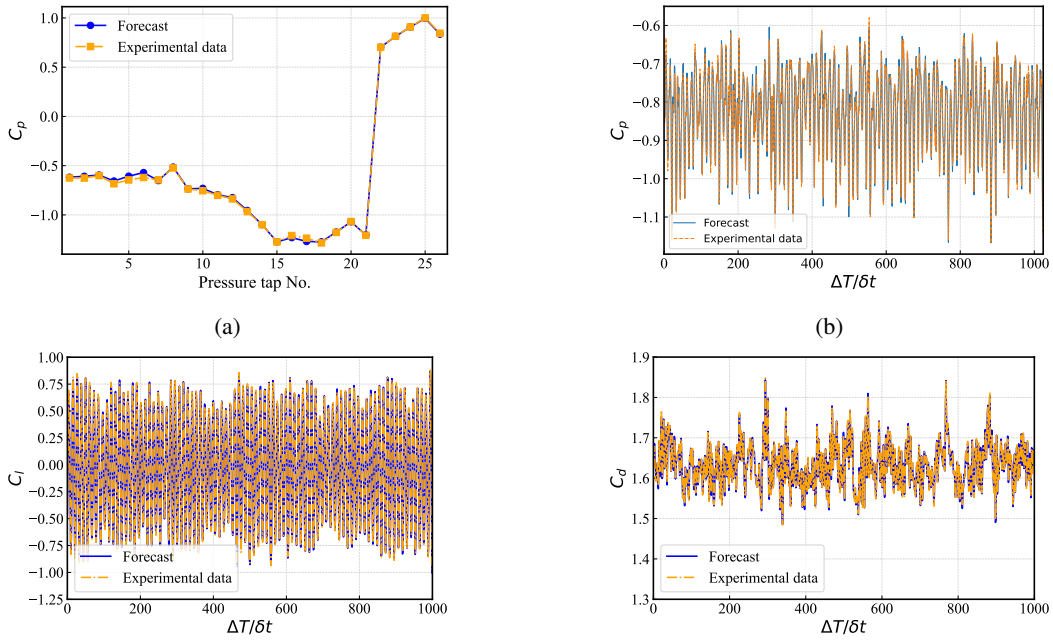

(a) (b)

(c) (d)

Figure 2: Physical information comparison between forecast results and experimental ground truth: (a)Instantaneous wall pressure distribution on rectangular cylinder walls, (b) Time history pressure coefficient with ΔT for the pressure tap 8 on the rectangular cylinder wall, (c) Time-series instantaneous lift coefficient C_l, (b) Time-series instantaneous drag coefficient C_d.

In addition, we investigate the power spectrum density (PSD) of the time-series information on both the experimental ground truth and forecast results. Figure 3 presents the power spectrum density of the pressure tap 4 on the wall. It is found that the PSD of forecast results align with experimental ground truth on pressure tap 4. More results and evaluations, including instantaneous wall pressure taps at different time instant, time-series wall pressure at different pressure taps, and PSD comparison of various pressure taps are shown in Appendix D.

There are also some future directions for improving the work. On one side, this time-series forecast has the same output length as the input information, and it can consider increase the output length or decrease the input length to reduce the data dependency. On the other side, current work use full dimensional data to forecast time-series bluff body wall pressure. While in some practical engineering cases, the full dimensional wall pressure data are not available or accessible. We can test the extrapolation capability in the future.

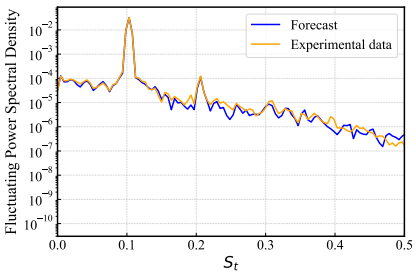

Figure 3: Fluctuation power spectrum density of time-series wall pressure information at pressure tap 4. The $x-$ axis is the Strouhal number $St = fU/D$, f is the vortex shedding dominate frequency, U is the approaching velocity, and D is bluff body width.

References

Z. Chen, Y. Liu, and H. Sun. Physics-informed learning of governing equations from scarce data. *Nature Communications*, 12(1):6136, 2021.

P. L. Fernández-Cabán, F. J. Masters, and B. M. Phillips. Predicting roof pressures on a low-rise structure from freestream turbulence using artificial neural networks. *Frontiers in Built Environment*, 4:68, 2018.

T. E. Ho, D. Surry, D. Morrish, and G. Kopp. The uwo contribution to the nist aerodynamic database for wind loads on low buildings: Part 1. archiving format and basic aerodynamic data. *Journal of Wind Engineering and Industrial Aerodynamics*, 93(1):1–30, 2005.

G. Hu, L. Liu, D. Tao, J. Song, K. 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, 2020a.

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

S. Li, X. Li, Y. Jiang, Q. Yang, M. Lin, L. Peng, and J. Yu. A novel frequency-domain physics-informed neural network for accurate prediction of 3D spatio-temporal wind fields in wind turbine applications. *Applied Energy*, 386:125526, 2025.

Y. Li, L. Yan, H. Gao, and G. Hu. A machine learning-augmented aerodynamic database of rectangular cylinders. *Physics of Fluids*, 36(7):075127, 07 2024. ISSN 1070-6631.

Z. Li, N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, and A. Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020.

Z. Li, D. Z. Huang, B. Liu, and A. Anandkumar. Fourier neural operator with learned deformations for pdes on general geometries. *Journal of Machine Learning Research*, 24(388):1–26, 2023.

J. Liu, K. Shum, T. 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.

J. Liu, K. T. Tse, and G. Hu. An aerodynamic database of synchronized surface pressure and flow field for 2D rectangular cylinders. *Engineering Structures*, 326:119506, 2025.

Z. Liu, Y. Wang, S. Vaidya, F. Ruehle, J. Halverson, M. Soljačić, T. Y. Hou, and M Tegmark. Kan: Kolmogorov-arnold networks. *arXiv preprint arXiv:2404.19756*, 2024c.

A. Mashhadi, A. Sohankar, and M. M. Alam. Flow over rectangular cylinder: Effects of cylinder aspect ratio and Reynolds number. *International Journal of Mechanical Sciences*, 195:106264, 2021.

J. H. Oh, G. A. Kopp, and D. R. Inculet. The uwo contribution to the nist aerodynamic database for wind loads on low buildings: Part 3. internal pressures. *Journal of wind engineering and industrial aerodynamics*, 95(8):755–779, 2007.

L. S. Pierre, G. Kopp, D. Surry, and T. Ho. The uwo contribution to the nist aerodynamic database for wind loads on low buildings: Part 2. comparison of data with wind load provisions. *Journal of Wind Engineering and Industrial Aerodynamics*, 93(1):31–59, 2005.

J. E. Santos, Z. R. Fox, A. Mohan, D. O'Malley, H. Viswanathan, and N. Lubbers. Development of the senseiver for efficient field reconstruction from sparse observations. *Nature Machine Intelligence*, 5(11):1317–1325, 2023.

K. Shimada and T. Ishihara. Application of a modified k–ε model to the prediction of aerodynamic characteristics of rectangular cross-section cylinders. *Journal of Fluids and Structures*, 16(4): 465–485, 2002.

A. Sohankar. Large eddy simulation of flow past rectangular-section cylinders: Side ratio effects. *Journal of Wind Engineering and Industrial Aerodynamics*, 96(5):640–655, May 2008. ISSN 0167-6105.

J. Tian, K. R. Gurley, M. T. Diaz, P. L. Fernandez-Caban, 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.

S. Wang, W. Cheng, R. Du, and Y. Wang. Unsteady RANS Modeling of Flow around Two-Dimensional Rectangular Cylinders with Different Side Ratios at Reynolds Number 6.85×10^5. *Mathematical Problems in Engineering*, 2020:2163928, September 2020. ISSN 1024-123X.

Y. Weng and S. G. Paal. Machine learning-based wind pressure prediction of low-rise non-isolated buildings. *Engineering Structures*, 258:114148, 2022.

A FKAN Model structure and physical frequency loss embeddement

The entire FKAN model is expressed as two main components including Fourier neural network and Kolmogorov–Arnold network. The entire model is shown in Figure 4.

We apply the fast Fourier transform (FFT) along the time dimension for time-series information $\mathbf{P}_i$ at i^{th} pressure tap on the bluff body wall:

$$\hat{\mathbf{X}}_{i,k} = \mathcal{F}(\mathbf{P}_i), \quad \hat{\mathbf{X}} \in \mathbb{C}^{C_{\text{in}} \times K}, \tag{7}$$

where i is the pressure tap index and k the frequency mode for each time-series pressure information, $\hat{\mathbf{X}}_{i,k}$ is the complex vector features of input wall pressure $\mathbf{P}_i$. For the first m low-frequency modes, we define a spectural convolutional network to learning complex weights $W_{i,k} \in \mathbb{C}$, giving

$$\hat{\mathbf{Y}}_{o,k} = \sum_{i=1}^{C_{\text{in}}} \hat{\mathbf{X}}_{i,k}\, W_{i,k}, \quad k = 0, \ldots, m-1. \tag{8}$$

where $\hat{\mathbf{Y}}_{o,k}$ is the complex vector features of the output wall pressure $\mathbf{P}_{o,i}$ at the i^{th} pressure tap. The time-series wall pressure output is reconstructed via the inverse Fourier transform:

$$\mathbf{P}_{o,i} = \mathcal{F}_t^{-1}(\hat{\mathbf{Y}}). \tag{9}$$

The above Eq. (7), (15), and (9) forms the first component as Fourier neural network (FNN).

For a random pressure tap i^{th} on the wall, we define a univariate learnable function to approximate the time-series wall pressure information using Fourier basis and polynomial terms:

$$\phi_i(x) = \sum_{k=1}^{K_{\max}} \Big(a_{i,k} \sin(kx) + b_{i,k} \cos(kx) \Big) + c_{1,i}\, x + c_{2,i}\, x^2 + b_i. \tag{10}$$

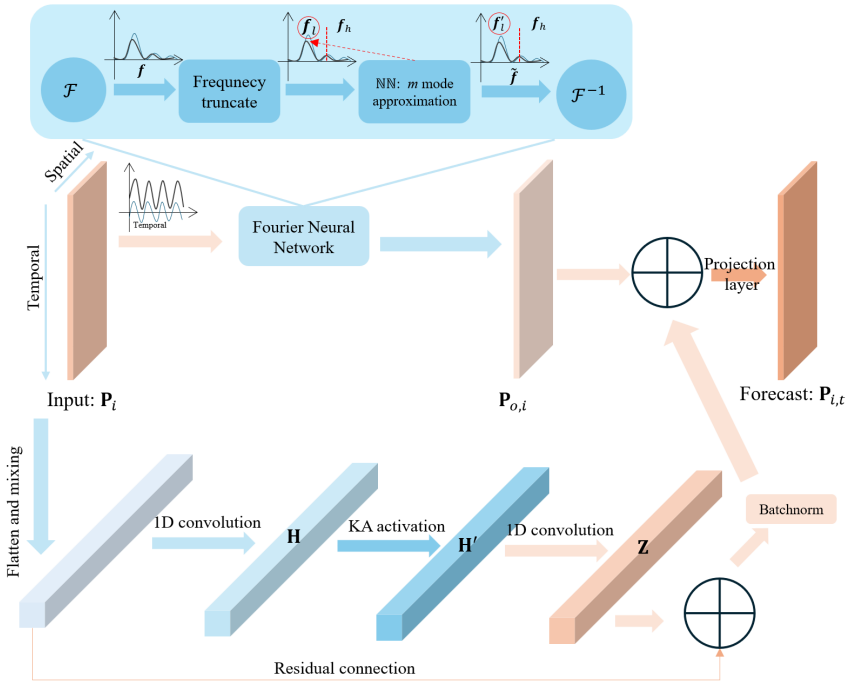

Figure 4: FKAN layer architecture details: in the above part, it is Fourier neural network, and the lower part is KAN.

This functional Kolmogorov–Arnold activation is applied elementwise on each pressure tap time series information:

$$Y_i = \phi_i(\mathbf{P}_i)\,, \tag{11}$$

where $\mathbf{P}_i \in \mathbb{R}^T$ is the input wall pressure information, and $\mathbf{Y}_i \in \mathbb{R}^T$ is the transformed time-series wall pressure information. The above Eq. (10) and (11) formulates the KAN layer in the KAN model.

Later, we define the KANBlock. The KANBlock performs channel mixing and nonlinearity on the entire pressure input $\mathbf{P} \in \mathbb{R}^{T\times N}$, where $T = 1024$ is the input time length, and $N = 26$ is the pressure taps number:

$$\mathbf{H} = \text{Conv1d}(\mathbf{P}), \tag{12}$$

After mixing, we perform a Kolmogorov–Arnold (KA) activation on i^{th} pressure tap as follows:

$$\mathbf{H}'_i = \phi_i(\mathbf{H}_i), \tag{13}$$

where ϕ is the Kolmogorov-Arnold actionation defined in Eq. (10). Based on the collected $\mathbf{H}'_i$, we form a $\mathbf{H}'$ for the entire pressure information $\mathbf{P}$, and perform a dimensional convolution based on the formulated $\mathbf{H}'$ as follows:

$$\mathbf{Z} = \text{Conv1d}(\mathbf{H}'). \tag{14}$$

We then apply the residual connection to $\mathbf{Z}$ and make a normalization:

$$\mathbf{Y} = \text{BN}(\mathbf{Z} + \mathbf{P}). \tag{15}$$

The KANBlock intergrates the above Eq. (12), (13), (14), and (15).

Each FKAN layer combines Fourier neural network, KANBlock, and skip connection as expressed below:

$$\mathbf{Y}_o = \text{ReLU}\Big(\text{BN}\big(\mathbf{FNN}(\mathbf{X}) + \text{KANBlock}(\mathbf{X}) + \text{Conv1d}_{1\times1}(\mathbf{X})\big)\Big), \tag{16}$$

Based on the FAKN layer, we use 4 FKAN layers to make the wall pressure forecast task.

The end-to-end model consists of the following stages:

$$\mathbf{P}_{l+1} = \text{FKAN}_l(\mathbf{P}_i), \quad l = 0, 1, \ldots, L-1 \tag{17}$$

where l is the number of FKAN layers, i.e., depth.

$$\mathbf{Y} = \text{Conv1d}_{1\times1}\{\sigma\big(\text{Conv1d}_{1\times1}(\mathbf{P}_l)\big)\}, \tag{18}$$

where $\sigma(\cdot)$ is a non-linear activation (i.e., GELU).

$$\mathbf{Y} \in \mathbb{R}^{T\times C_{\text{out}}}. \tag{19}$$

B Experimental Validation and Data Processing

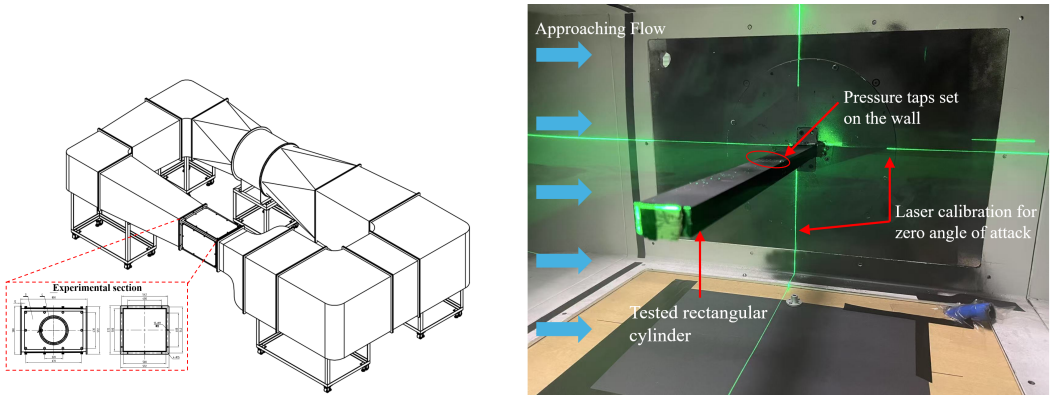

(a) (b)

Figure 5: Wind tunnel testing setups for wall pressure collection of the bluff body: (a) Closed-loop wind tunnel for the wind tunnel experiment, (b) Experimental setups for wind tunnel testing. Green lines mean the calibration for the zero angle of attack.

The typical bluff body, a rectangular cylinder of side ratio = 1.5 under zero angle of attack, is tested in the closed-loop wind tunnel as shown in Figure 5a. To ensure the approaching flow to the rectangular cylinder strictly follows zero angle of attack, we calibrated the rectangular cylinder arrangement with a 12-line laser calibration instrument before the experiment shown in Figure 5b. The pressure scanners used in the experiment to sample wall pressure information is *Scanivalve MPS4262* pressure scanners. In this experiment, the experimental surrounding condition is around 26 degrees Celsius, and the experimental temperature remains the best working condition for the pressure scanning system [Liu et al., 2025]. The length of the pressure tube is less than 0.6 *m* to reduce the distortion of wall pressure information in the experiment. The pressure scanning starts working after the approaching flow is consistent and the turbulent flow has been fully developed in the rectangular cylinder wake. The rectangular width D is 30 mm, and length (L) is 45 mm. On all four walls, the distance between each pressure tap is 5 mm shown in Figure 6. The pressure sampled by the pressure scanners is a relative pressure, indicating that the pressure obtained has been subtracted by the reference pressure p_{ref}. With the collected pressure information on the wall of the rectangular cylinder, the pressure information is processed to validate the precision of the experimental data by comparing the collected information with existing research data. For flow past a rectangular cylinder under the angle of attack 0°, the Strouhal number (St), the mean drag coefficient ($\overline{C}_d$), and the standard deviation value of the lift coefficient ($C_{l,std}$) determine the precision of the pressure information. Thus, these aerodynamic statistical features are calculated to validate the experimental data collected from the wind tunnel test. The mean drag coefficient $\overline{C}_d$ is calculated as follows:

$$\overline{C}_d = \frac{1}{N}\sum_{i=1}^{N} C_{d,i}, \tag{20}$$

where $N = 24,000$ is the total number of snapshots collected from the wind tunnel experiment. $C_{d,i}$ is the instantaneous drag coefficient, and the instantaneous drag coefficient is calculated by the

instantaneous wall pressure difference on both leading and trailing walls of the rectangular cylinder as follows:

$$C_{d,i} = \frac{1}{n_d} \sum_{j=1}^{n_d} \frac{p_{j,r}^i - p_{j,l}^i}{1/2\rho U^2}, \tag{21}$$

where $n_d = 5$ is the number of pressure taps on the leading and trailing walls for the rectangular cylinder of the side ratio 1.5 in this study. $U = 10\ m/s$ is the approaching wind velocity in this experiment. Here, $\rho = 1.225\ kg/m^3$is the air density. $p_{j,l}^i$ and $p_{j,r}^i$ respectively denote the wall pressure information sampled by the jth spatial pressure scanners on leading and trailing walls at the particular temporal instant i. Here, i ranges in 24,000, which is the entire temporal domain. It should be noted that the instantaneous darg coefficient ($C_{d,i}$) here is the pressure difference of the leading and trailing walls, without considering the friction drag under this condition, as when the Reynolds number is greater than 200, the friction drag can be neglected for rectangular cylinder at $SR = 1.5$ [Mashhadi et al., 2021]. The instantaneous lift coefficient is calculated as follows:

$$C_{l,i} = \frac{1}{n_l} \sum_{j=1}^{n_l} \frac{p_{j,low}^i - p_{i,up}^i}{1/2\rho U^2}, \tag{22}$$

where $n_l = 8$ is the number of pressure taps on the upper and lower walls, $p_{j,low}^i$ and $p_{i,up}^i$ denote the wall pressure information sampled by the pressure tap at the jth spatial pressure tap on the upper and lower wall at the particular time instant i, and i ranges within 24,000. After getting the instantaneous lift coefficient $C_{l,i}$, the standard deviation of the lift coefficient can be calculated based on the temporal domain as follows:

$$C_{l,std} = \sqrt{\frac{1}{N} \sum_{i=1}^{N} (C_{l,i} - \overline{C}_l)^2}, \tag{23}$$

where $N = 24,000$ is the total number of snapshots as in Eq. (20), $\overline{C}_l$ is the temporal average of the lift coefficient $C_{l,i}$. Besides, the fast Fourier transform (FFT) is performed based on the temporal history of lift coefficient to obtain the wake vortex shedding dominant frequency f, which is used to calculate the Strouhal number (St). The Strouhal number is calculated as follows:

$$St = \frac{fD}{U}, \tag{24}$$

where $D = 0.03$ m is the characteristic length and $U = 10$ m/s is the approaching velocity. Based on the wall pressure information collected, the dominant vortex shedding frequency is calculated as around 34.3 Hz. The pressure scanning frequency of 400 Hz can be sufficient to capture fluctuating features. Moreover, the Reynolds number is defined as

$$Re = \frac{UD}{\nu}, \tag{25}$$

where ν=1.51$\times 10^{-5}$ m^2/s is the kinematic viscosity of the air. The above three aerodynamic characteristics, including St, $\overline{C}_d$, and $C_{l,std}$, can be compared with historical research data, as shown in Table 1. When comparing aerodynamic statistics with those in existing literature, it can be seen that the wall pressure information collected from this wind tunnel experiment is in good agreement with existing research data [Shimada and Ishihara, 2002, Sohankar, 2008, Wang et al., 2020]. Statistical verification indicates that the wall pressure measurements can serve as a basis for further model training and testing in the following sections.

In the data sampling for FKAN model training, we follow the time sequence to get time-series snapshot of instantaneous wall pressure information shown in Figure 7. The length of the collected dataset is 24,000, and we sample 1024 time-series snapshots as input and the following time-series 1024 snapshots as the labeled forecast results to train the model. In total, there are 21,952 pairs of labeled data for the input and forecast. We divide 30% of the dataset as the train dataset of 6586 pairs, and the other 70% of the dataset is used as the test dataset of 15,366 pairs.

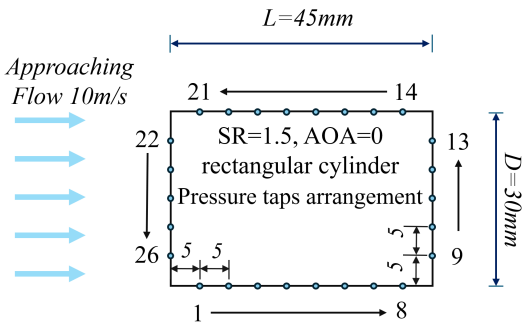

Figure 6: Tested bluff body dimension information and pressure taps arrangement.

Table 1: Aerodynamic statistics (St, $\overline{C}_d$, $C_{l,std}$) comparison between wind tunnel experimental data here and existing research data for the rectangular cylinder of side ratio 1.5 under the angle of attack 0°.

Literature	Re	St	$\overline{C}_d$	$C_{l,std}$
Shimada and Ishihara [2002]	24,000	0.103	1.62	0.50
Sohankar [2008]	100,000	0.095	1.62	/
Wang et al. [2020]	685,000	/	1.63	0.67
This work	around 20,000	0.103	1.64	0.51

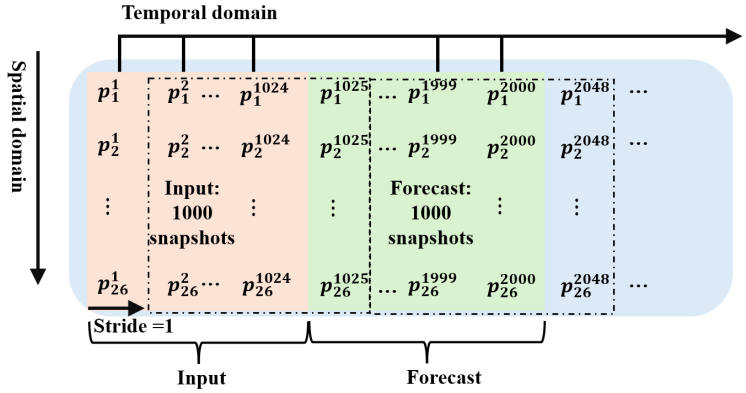

Figure 7: Dataset sampling for train and test data with sampling window:

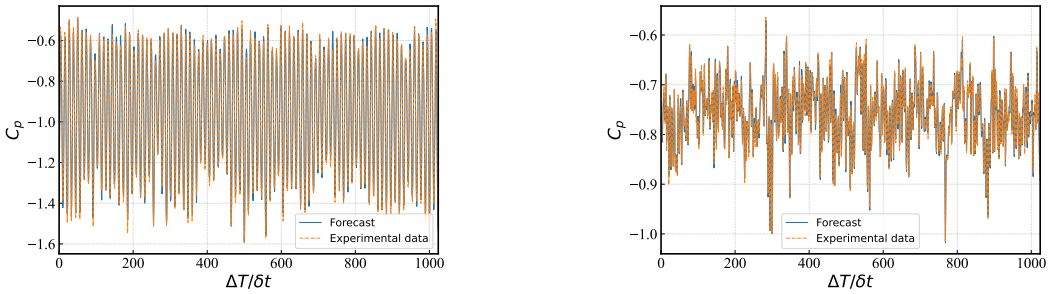

(a) (b)

Figure 8: Time-series wall pressure information comparison between forecast results and experimental ground truth at different pressure taps: (a)Pressure tap 4, (b) Pressure tap 10

C Model training settings and device usage

In the model training, we use *Nvidia A40 GPU* for model training, and the memory used for this light FKAN model is around 1.6 GB. During model training, we set the hyperparameters for the FKAN as shown in Table 2.

Table 2: Hyperparameters for FKAN model training.

Hyperparameters	Values
Learning rate	0.0001
Batch size	16
$epochs$ in Eq. (2)	3000
β_{end} in Eq.(3)	0.001
FKAN layer depth l in Eq (17)	4
m modes of FNN in Eq (8)	64

D More results evaluation and comparison

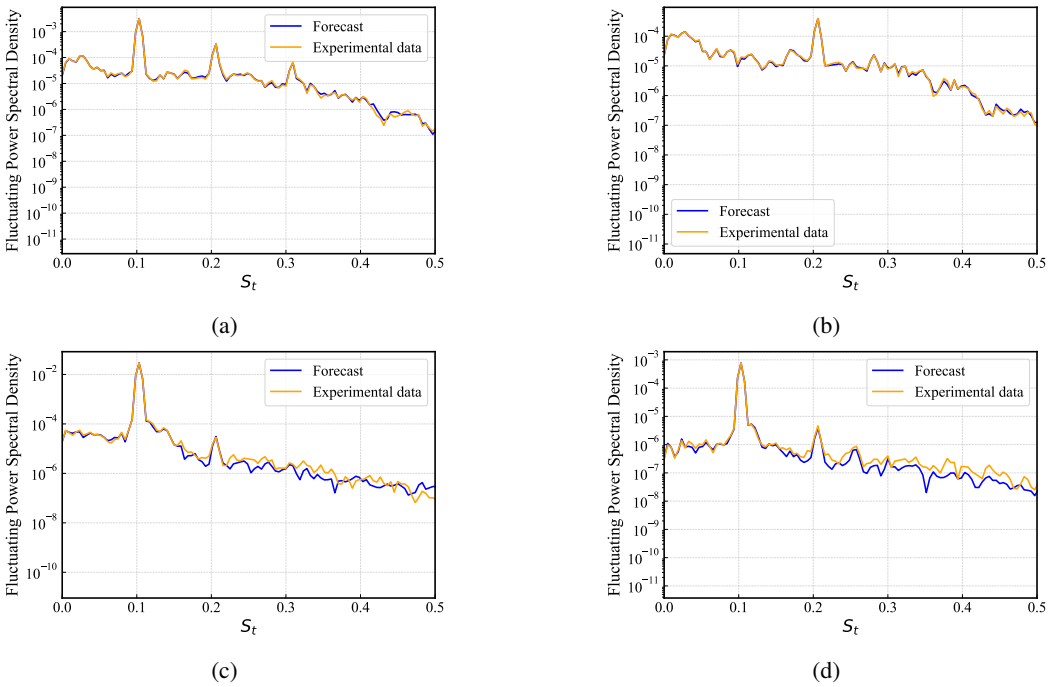

Figure 9: Power spectrum density for time-series wall pressure information at different pressure taps: (a) Pressure tap 8, (b)Pressure tap 10, (c) Pressure tap 17, (d) Pressure tap 21.

Figure 8 presents more results on time-series wall pressure information at different pressure taps, including pressure tap Nos. 4, and 10. From the time-series comparison between experimental ground truth and forecast results, it is found the FKAN model can capture the dynamic features at different pressure taps. Figure 9 show the fluctuating power spectrum density for time-series information at different pressure taps to estimate the physical interruption of FKAN model. From the comparison, it can be seen that FKAN model can accurately capture the frequency peaks, including dominant frequency $St \approx 0.103$ shown in Figure 9a, 9c, 9d, twice of the dominant frequency ($2St \approx 0.21$) shown in all four figures of Figure 9, and even three times of the dominant frequency ($3St \approx 0.31$). Figure 10 demonstrate the instantaneous wall pressure distribution at various time instant indices.

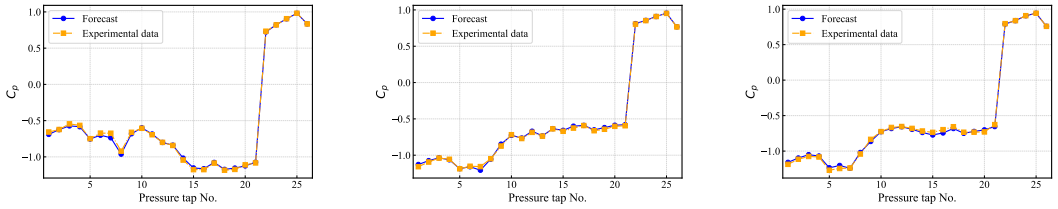

(a) (b) (c)

Figure 10: Various instantaneous snapshot of wall pressure distribution at different time instant index: (a) $\Delta T/\delta t = 83$, (b) $\Delta T/\delta t = 383$, (c) $\Delta T/\delta t = 583$

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: The abstract and introduction clearly outline the contributions and scope of the paper.

   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: The paper discuss the limitations of the work performed.

   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: The paper does not include theoretical results.

   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: The paper fully discloses all necessary information required to reproduce the main experimental results, adequately supporting the main claims and conclusions.

   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: [Yes]

   Justification: The data and code, along with all relevant information, will be available upon paper 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, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

   Answer: [Yes]

   Justification: The paper specifies all necessary training and test details required to understand the results clearly in the code.

   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: The paper does not report error bars or other appropriate information regarding the statistical significance of the experiments in a suitable and correctly defined manner.

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: [Yes]

   Justification: The paper provides sufficient information regarding the computer resources required to reproduce the experiments.

   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: The research conducted in the paper fully conforms with the NeurIPS Code of Ethics in all respects.

   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: [No]

    Justification: The paper only discusses potential positive societal impacts of the work performed.

    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: The paper poses no such risks.

    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: [Yes]

Justification: The paper properly credits the creators or original owners of all assets used, and the licenses and terms of use are explicitly mentioned and respected throughout the manuscript.

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: The paper does not release new assets.

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: The paper does not involve crowdsourcing nor research with human. subjects.

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: The paper does not involve crowdsourcing nor research with human. subjects.

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: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

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.