

Comparing Data-Driven Approaches for Predicting the Flow-Induced Vibrations of Elliptical Cylinders

Jonathan C.C. Lo^{1*}, Mark C. Thompson¹, Kerry Hourigan¹ and Jisheng Zhao¹

¹ Fluids Laboratory for Aeronautical and Industrial Research (FLAIR), Department of Mechanical and Aerospace Engineering, Monash University, Victoria 3800, Australia

*mailto: jonathan.lo1@monash.edu.au

Abstract

This data-driven study seeks to predict the amplitude response of an elastically-mounted elliptical cylinder by leveraging previous experimental investigations to train machine learning algorithms. The elliptical cylinder used had an elliptical ratio of $\varepsilon = b/a = 5$, where a and b are the stream-wise and cross-flow dimensions, respectively. Specifically, we evaluated the performance of k -nearest neighbours, decision tree, gradient-boosted decision tree, and random forest models. Using the structural damping ratio, Reynolds number, and reduced velocity as inputs, each method was optimised using particle swarm optimisation with respect to the mean absolute error loss. The k -nearest neighbours had the lowest loss value and was hence chosen as the best algorithm, with performance on the test data indicating that it was able to capture the main trends of the amplitude response. This demonstrates the ability of machine learning to predict the effect of structural damping on the flow-induced vibration response of an elliptical cylinder, with potential applications in the field of renewable energy where generator loads can be modelled as applied damping.

1 Introduction

Flow-induced vibration (FIV), arising from the coupled interaction between a fluid and a structure, is an important phenomenon prevalent in various areas of engineering (Bearman, 1984; Naudascher & Rockwell, 2005; Williamson & Govardhan, 2004). Notably, FIV can lead to catastrophic structural failure or long-term fatigue, but it can also be potentially harnessed as a source of green energy (Bernitsas *et al.*, 1984; Lv *et al.*, 2021).

FIV of an aeroelastic bluff body can often be manifested as two body-oscillator phenomena: *vortex-induced vibration* (VIV) and galloping. Fundamentally, VIV is caused by the periodic or quasi-periodic shedding of vortices in a *vortex street* pattern, while galloping is a movement-induced instability characterised by a linear amplitude growth with flow speed past a critical value (Blevins, 1990). Depending on the flow conditions and geometric properties (e.g. the geometric shape, structural damping, and mass ratio), VIV and galloping may occur separately or simultaneously. As such, previous studies have focused on circular and square cylinders since FIV exclusively manifests as VIV and galloping, respectively, allowing both phenomena to be studied independently. However, much less research has been conducted on the coupling between both forms of FIV that can occur for symmetry-breaking geometries. Recent studies (see Nemes *et al.*, 2012; Zhao *et al.*, 2018c, 2014) have shown that rich, complex, and non-linear dynamics can also be observed when both FIV phenomena concurrently manifest. Although mixed-mode FIV occurs in many real-world scenarios, the coupling of VIV and galloping makes traditional quasi-steady theory inapplicable for the prediction and characterisation of the resultant dynamics.

One example of mixed-mode FIV relates to cylinders with elliptical geometries (Zhao *et al.*,



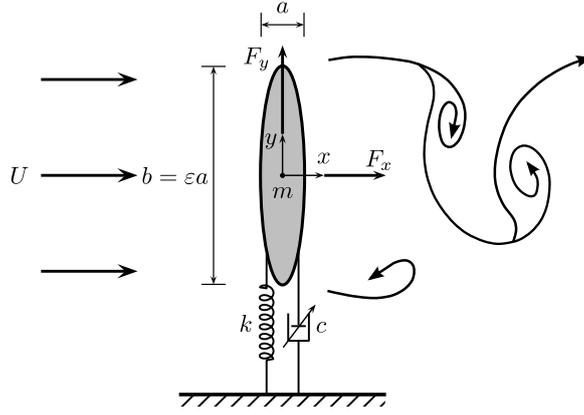


Figure 1. A schematic defining the problem of interest: an elastically mounted elliptical cylinder constrained to oscillate transverse (y) to the free stream flow of velocity U and in the positive x direction. Here, the geometry is characterised by the elliptical ratio $\varepsilon = b/a$, where a and b are the stream-wise and cross-flow dimensions, respectively. Additionally, m is the oscillating mass, k denotes the spring constant, c is the adjustable structural damping, and F_x and F_y represent the respective drag and the transverse (lift) fluid forces acting on the body, respectively.

2018b). Since the circular cylinder is considered a special case of the elliptical geometry whose FIV response is only VIV due to the inherent rotational symmetry of the system, introducing deformation through the eccentricity of the cross-section causes the profile to be susceptible to movement-induced instability. A case in point is our recent and unpublished discovery of large-scale oscillations (occurring in a regime coined the *hyper-branch* (Zhao *et al.*, 2023)) for an elastically mounted elliptical cylinder submerged in water. This system is modelled as a linear second-order damped harmonic oscillator of form

$$m\ddot{y}(t) + c\dot{y}(t) + ky(t) = F_y(t), \quad (1)$$

which relates the total fluid force acting in the cross-flow direction ($F_y(t)$) to the system mass (m), spring constant (k), structural damping constant (c), and cylinder displacement ($y(t)$). Figure 1 shows the schematic of the cylinder where eccentricity of the elliptical geometry is defined by the elliptical ratio, $\varepsilon = b/a = 5$. With observed amplitudes nearly eight times the cross-flow dimensions of the cylinder, Lo *et al.* (2023) then demonstrated that the hyper-branch oscillations can be dampened and eventually suppressed with increasing structural damping, as shown by the amplitude response in figure 2. The amplitude (A), flow speed (U), and structural damping (c) are represented by their non-dimensional forms: the normalised amplitude parameter $A^* = A/b$, reduced velocity $U^* = U/(f_{nw}b)$, and the structural damping ratio accounting for added mass effects $\zeta = c/(2\sqrt{k(m+m_A)})$. Here, f_{nw} and m_A denote the natural frequency of the elliptical cylinder when submerged in quiescent water and potential added mass, respectively.

However, understanding the effect that parameters (e.g. ζ and U^*) have on the amplitude response of elastically mounted elliptical cylinders using first principles and equation-based models is made difficult by the multi-scale and non-linear nature of the FIV. To address this issue, data from previous studies could be leveraged by machine learning algorithms to gain further insights into the amplitude response arising from FIVs. Although the review by Brunton *et al.* (2020) has shown the success of machine learning in many fluid dynamics applications, few studies have applied this data-driven approach to FIV research (Lin *et al.*, 2021; Ma *et al.*, 2020; Mei *et al.*, 2021) and none in predicting, from experimental parameters, the hyper-branch response of elliptical cylinder FIV.

This study presents a comparison of machine learning models for the amplitude response prediction of an $\varepsilon = 5$ elliptical cylinder based on the applied structural damping of the system. The article proceeds with the methodology in § 2, providing an overview of the data and models used. The performance of the trained and optimised models were then discussed in § 3, with the conclusions of this study drawn in § 4.

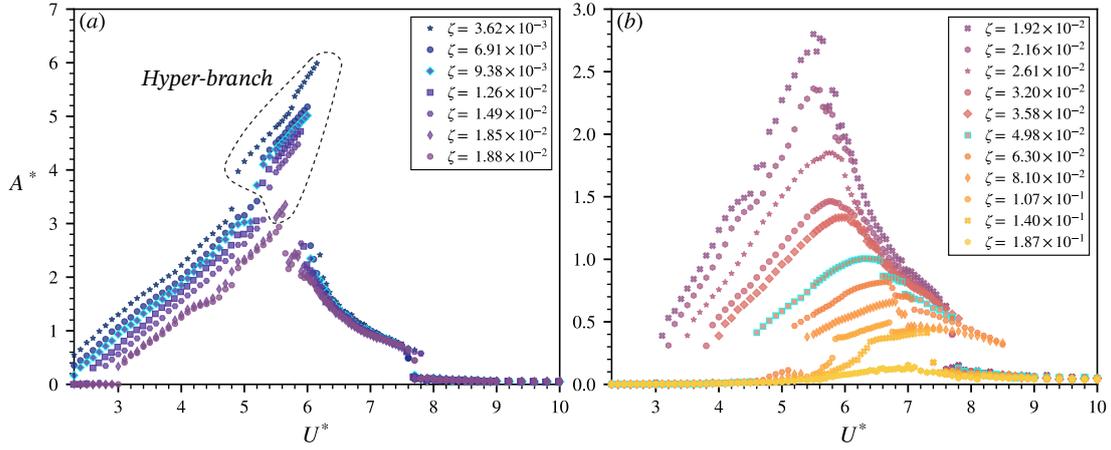


Figure 2. Amplitude response (A^*) of the elliptical cylinder as a function of reduced velocity (U^*) for various applied structural damping (ζ) values over a Reynolds number range of $980 \leq \text{Re} \leq 4410$. Using decreasing U^* increments, the amplitude responses corresponding to the presence and suppression of the hyper-branch are shown in (a) and (b), respectively. Markers with cyan edges indicate data points that have been set aside as the test data.

2 Methodology

2.1 Data Collection for Model Training

Figure 2 shows our unpublished dataset (Lo *et al.*, 2023) on which the models will be trained, optimised and tested. The experiments were conducted within the free-surface recirculating water channel of the Fluids Laboratory for Aeronautical and Industrial Research (FLAIR) at Monash University. To implement free transverse vibration, the elliptical cylinder was mounted on a low-friction air-bearing rig (Zhao *et al.*, 2018c) and was elastically constrained by high-precision extension springs. The structural damping of the air-bearing rig was controlled using an electromagnetic damper system (Soti *et al.*, 2018).

The preliminary results demonstrated that the FIV response of the elliptical cylinder as a function of U^* , both in shape and in amplitude, is dependent on ζ . The body vibration was observed to be strongly associated with a fluid-structure synchronisation, where the vortex shedding frequency “locks” onto the body oscillation frequency. For the hyper-branch, additional higher order harmonic components in the lift force were also observed, suggesting that the large-scale oscillations were associated with a combined effect of VIV and galloping. It was found that increasing ζ reduces the magnitude of the body vibration, with the hyper-branch suppressed for $\zeta \geq 1.92 \times 10^{-2}$. The presence of hysteresis was also observed, with the FIV response depending on the direction of U^* increments used to construct the response curves. Due to the water channel limits, measurements could not be taken at all U^* values of interest for increasing U^* and, as such, the focus of this study is only on data from decreasing U^* increments.

Eighteen damping values were tested in order to investigate the effect of damping on the observed hyper-branch phenomenon, with a total of 1613 points in the dataset. To have an unbiased and accurate measurement of model accuracy after training and optimisation, two damping values were taken as the 182-point test data set. Shown by the cyan marker edge in figure 2, $\zeta = 9.38 \times 10^{-3}$ and 4.98×10^{-2} were chosen as they respectively represent the two main FIV responses — with and without the hyper-branch. As such, the training/testing split is approximately 90-10%. In this study, training data are used for training (§ 2.3), model optimisation (§ 2.4), and model selection. The test set is withheld and only used to provide an unbiased evaluation of the accuracy of the final model at the end of the study.

Three experimental parameters are chosen as the inputs for the machine learning algorithms of

interest: structural damping ratio (ζ), reduced velocity (U^*), and Reynolds number ($Re = (Ub)/\nu$, where ν is the kinematic viscosity of the water). Although the aim of the study is to predict the amplitude response as a function of reduced velocity based on a given damping ratio, the Reynolds number was also included to account for variations in the natural frequency ($f_{n,w}$) that arises from different damping.

2.2 Machine Learning Implementation

Model	Optimised Hyper-parameters	Lowest Loss Value
<i>k</i> -Nearest Neighbours	Algorithm: Ball Tree, Weight: Distance, Leaf size: 84, Neighbours: 2, Metric power: 1	$(1.06 \pm 0.02) \times 10^{-1}$
Decision Tree	Tree Depth: 27, Min. samples before split: 4, Min. samples in leaf node: 1	$(1.29 \pm 0.02) \times 10^{-1}$
Random Forest	Tree Depth: 56, Min. samples before split: 5, Min. samples in leaf node: 1, Trees: 156	$(1.30 \pm 0.02) \times 10^{-1}$
Gradient-Boosting Decision Tree	Tree Depth: 37, Min. samples before split: 6, Min. samples in leaf node: 4, Trees: 41, Learning Rate: 1.622e-1	$(1.22 \pm 0.02) \times 10^{-1}$

Table 1. The hyper-parameter values corresponding to the lowest mean absolute value loss for each machine learning model were obtained via particle swarm optimisation. The loss for each model was estimated as the mean of 8-fold cross validation repeated 100 times, and presented along with their 95% confidence intervals.

All machine learning components of this study were implemented in Python, with the Scikit-Learn library utilised to build the models. Chosen due to their popularity and robustness in many applications, table 1 shows the architectures tested in this study for the amplitude prediction of elastically mounted elliptical cylinders undergoing FIV. In contrast with popular deep machine learning methods such as neural networks, all the architectures tested in this study are considered to be shallow learners due to their limited representation spaces. The amplitude response predictions made by the models utilise the existing experimental data to find relationships between the inputs and outputs. Aside from the *k*-nearest neighbours algorithm which performs regression directly from searching the data and taking the majority vote of *k* nearest data points, the input-output relationship for the other models are determined by the model parameters that are learned during training.

The principle of this process is to find the aforementioned parameters that maximise the model accuracy through the minimisation of a predefined function known as the *loss*. Given the existence of instantaneous amplitude transitions in the data, the mean absolute error (MAE) was chosen as the loss function over the more popular mean-squared error (MSE). This is due to the latter's bias towards accurately fitting the amplitude jumps at the expense of the entire response curve. To compare the performance between models, *k*-fold cross-validation is often used to provide an estimate of loss. For this approach, the training data are split into *k* sets, where one set is selected for validation and is used to evaluate the performance of a model trained on the remaining *k* - 1 sets. This process is repeated *k* times such that each set is the validation set once, with the resultant *k* performance scores averaged to provide the final estimate. As the purpose of this study is predict the amplitude response based on the structural damping, 8-fold cross validation was utilised with the splits conducted across the ζ axis of the $\{\zeta, U^*, Re\}$ input space.

Beyond the aforementioned parameters of the machine learning algorithms, hyper-parameters also have an effect on model performance. Defined as a parameter whose value is used to control the learning process, hyper-parameters are model-dependent and must be optimised outside of the training process. § 2.3 will briefly discuss the algorithms investigated in this study, while further details on how their hyper-parameters were optimised can be found in § 2.4.

2.3 Machine Learning Models

This subsection discusses each of the four machine learning methods tested in this study. It should be noted that with the exception of the k -nearest neighbours model, the other algorithms utilise the decision tree as the base estimator.

The method of **k-nearest neighbours** is an algorithm that performs regression by interpolating between the labels of the k nearest samples of the training data to a given input. The hyper-parameters that govern this prediction process include the number of neighbours k , the Minkowski metric power parameter, the weights (either uniform or weighted by the inverse distance from sample to input), and the algorithm and leaf size used to determine the nearest points (either KD or Ball tree).

The **decision tree** is constructed using the Classification And Regression Trees (CART) algorithm to recursively and greedily split the data (predicated on their location in the input parameter space) with the aim of minimising the loss at each split. The total data initially forms a single set (*root node*), with each rule-based decision splitting the node into subsequent child nodes that are then further split as part of a recursive process. The last nodes of the decision tree are known as *leaf* nodes, and are used by the model to make the final prediction. The hyper-parameters considered in this study for the decision tree include the maximum depth of the tree, the minimum number of samples required to be at a leaf node, and the minimum number of samples required to split an internal node.

Gradient-boosted decision trees are an ensemble technique built upon numerous weak learners (i.e. decision trees) as the base estimator. The first weak learner is trained directly on the data whilst subsequent trees are trained to predict the residual loss for each sample of the previous trees. While increasing the number of learners generally reduces the model error, this may introduce over-fitting. As such, along with the hyper-parameters of the decision tree as the base learner, additional hyper-parameters considered for this algorithm include the number of learners and the learning rate. The learning rate minimises over-fitting by reducing the contribution of each subsequent learner to the overall prediction.

The **random forest** is another ensemble technique also utilising the decision tree as the weak learner, with the random output taken as an average of the multiple learners. To improve accuracy and reduce over-fitting, bootstrapping aggregating (also known as bagging) is used to ensure that each learner is trained on a random subset of the training data. Here, the bagging process involves drawing n (total number of observations in the data set) samples from the dataset *with* replacement to train each tree. Along with the hyper-parameters of the base estimator in the decision tree learner, the number of learners is an additional hyper-parameter of the random forest.

2.4 Hyper-parameter Tuning with Particle Swarm Optimisation

To maximise the accuracy of each algorithm tested in § 2.3, their hyper-parameters must be carefully chosen to optimise the performance for a given data set. Often, the model is trained multiple times in a grid search through the hyper-parameter space to find the combination that yielded the best performance as measured by the minimisation of the loss. However, with the success by Lin *et al.* (2021) in utilising the particle swarm optimisation to fine-tune their machine learning models, a similar process will also be used to optimise the architecture of each model in this study.

Particle swarm optimisation is an algorithm belonging to a sub-field of artificial intelligence known as swarm intelligence. Defined as “*the emergent collective intelligence of groups of simple agents*” by Bonabeau *et al.* (1999), swarm intelligence is biologically inspired and mimics the self-organisation and division of labour that is commonly found in social creatures (e.g. ants, bees, bats, etc.). Particle swarm optimisation will be used in this study due to its popularity (Slowik & Kwasnicka, 2017) and its success by Lin *et al.* (2021) in optimising their machine learning models. Inspired by the behaviour of bird flocks, particle $j = 1, 2, \dots, J$ at the $(i + 1)^{\text{th}}$ iteration of particle swarm optimisation with a total of J agents is mathematically defined as (Shi & Eberhart, 1998)

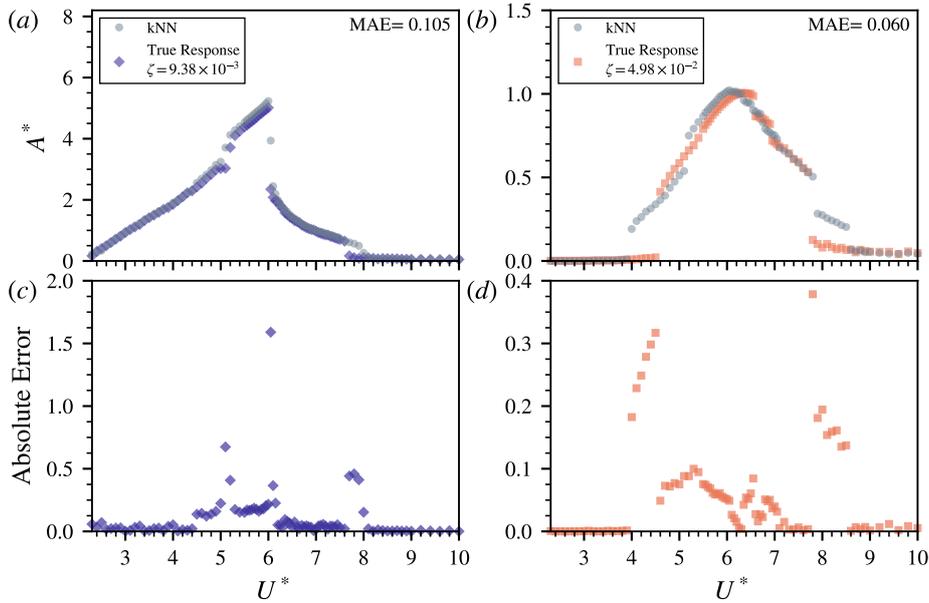


Figure 3. (a-b) The amplitude response curves predicted by k -nearest neighbours for $\zeta = 9.38 \times 10^{-3}$ and $\zeta = 4.98 \times 10^{-2}$ in the test data set. The MAE for both cases are shown in the top left of the both plots, with the absolute error between the true and predicted response as a function of U^* presented in (c-d).

$$\vec{v}_j^{i+1} = \omega \vec{v}_j^i + c_1 r_1 (\vec{P}_{\text{best},j}^i - \vec{x}_j^i) + c_2 r_2 (\vec{G}_{\text{best}}^i - \vec{x}_j^i) \quad (2)$$

$$\vec{x}_j^{i+1} = \vec{x}_j^i + \vec{v}_j^{i+1}, \quad (3)$$

where \vec{x}_j^i and \vec{v}_j^i is the position and velocity of particle, respectively, with \vec{x}_0 and \vec{v}_0 initialised as random values in the multi-dimensional parameter space. Here, r_1 and r_2 are outputs of an uniform random number generator between 0 and 1, introducing randomness into the system. $\vec{P}_{\text{best},j}^i$ and \vec{G}_{best}^i are the personal best position of particle j up to the current iteration and the global best position of the entire swarm, respectively, as measured by the function the algorithm seeks to minimise. The hyper-parameters of this algorithm are the inertia factor ω and the acceleration constants c_1 and c_2 , which affect how the particles explore the solution space and the convergence of the algorithm (Shi & Eberhart, 2001).

For this study, particle swarm optimisation was implemented via the Pyswarms library using 200 particles evolving over 200 iterations. The hyper-parameters were chosen to be $\{\omega = 0.729, c_1 = c_2 = 1.49445\}$, in line with the values suggested by Shi & Eberhart (2001) for better convergence. The function to minimise in this study is the MAE loss as estimated by 8-fold cross validation. To account for uncertainty in our estimate with 95% confidence intervals, the loss of the optimised models reported in table 1 were taken as the mean of 100 repeated 8-fold cross validations with randomly generated folds. However, by assuming that the ideal hyper-parameters for a given model is approximately invariant to fold choice, computing time is decreased during particle swarm optimisation by using a single 8-fold cross validation with constant folds for each ‘‘particle’’ evaluation.

3 Results

Table 1 shows the optimal hyper-parameters and the corresponding lowest MAE (estimated using 8-fold cross validation repeated 100 times) achieved for each model as identified with particle swarm optimisation. Accounting for the 95% confidence interval associated with each estimate, only the

difference in loss between decision tree and random forest (whose confidence intervals overlapped) was not statistically significant. Interestingly, the k -nearest neighbours was the simplest yet the best-performing out of all the algorithms tested with the lowest MAE estimate of 1.06×10^{-1} . This conclusion was in contrast to that drawn by Lin *et al.* (2021) who showed that the gradient-boosted decision tree was better than the random forest, decision tree, and k -nearest neighbours models. This difference could be due to their larger input parameter space of four variables as compared to that investigated by our current study (three input variables of which two, the Reynolds number and reduced velocity, are closely related). Additionally, the nature of the data and the use of a different loss (MSE in the case of Lin *et al.* (2021)) might also affect the accuracy of each model differently.

As such, the optimised k -nearest neighbours model was hence taken as the final chosen model and evaluated on the previously-unseen test data set to obtain the generalised error. Figures 3(a-b) shows the resultant k -nearest neighbours-predicted amplitude responses, while figure 3(c-d) shows the absolute error as a function of U^* . For the hyper-branch test case ($\zeta = 9.38 \times 10^{-3}$), the response curve was well predicted as shown by the MAE of 0.105. However, the main sources of error were found between $U^* \in [5, 6]$ and $U^* \approx 8$, which corresponds respectively to underestimating the hyper-branch regime and overestimating the U^* at which the step-wise transition to negligible amplitudes occur. For the second test case, where the hyper-branch is suppressed ($\zeta = 4.98 \times 10^{-2}$), the k -nearest neighbours was able to predict the main trends of the curve as demonstrated by the MAE of 0.060, but finer features such as the amplitude transitions between $U^* \in [6.5, 7]$ were lost. Additionally, the k -nearest neighbours also incorrectly predicted step-wise amplitude transitions in the response curves, which are non-existent in the true response. The cause of these errors may be due to the nature of the prediction mainly being an interpolation between the two adjacent damping ratios ($\zeta = 3.58 \times 10^{-2}$ and 6.3×10^{-2}) and does not include knowledge from other parts of the data set.

To improve the results obtained in this study, future work may include using amplitude responses of more structural damping ratios as part of the training data which would help increase model accuracy. In addition, extending the study to also include data from increasing U^* increments will allow the resultant models to account for the role that hysteresis plays in the amplitude response of elliptical cylinder FIV. Finally, as the models tested in this study are considered “shallow” due to the small representation spaces, including deep learning algorithms (such as neural networks) for comparison in the future would elucidate whether more complex models can improve predictive accuracy.

4 Conclusions

The performance of four different machine learning algorithms for the prediction of the amplitude response pertaining to $\varepsilon = 5$ elliptical cylinder FIV has been investigated. Utilising the damping ratio, reduced velocity and Reynolds number as the input parameters and optimising each tested architecture using particle swarm optimisation, we found that the lowest MAE loss (estimated using 8-fold cross validation) was achieved by the k -nearest neighbours model. As this performance over the other algorithms was considered statistically significant when accounting for 95% confidence intervals, the k -nearest neighbours was chosen as the final model. When evaluated on the test data, the MAE was 0.105 and 0.060 for $\zeta = 9.38 \times 10^{-3}$ and 4.98×10^{-2} , respectively. Since structural damping is an analogue for the loads applied by a generator, it is hoped that a robust approach to predicting the effects of structural damping on the FIV amplitude response of elliptical cylinders will assist in the development of devices that employ this geometry for efficient renewable energy generation.

Acknowledgements

J.C.C.L. acknowledges the financial support of an Australian Government Research Training Program Scholarship. We would also like to acknowledge the financial support of the Australian Research

Council (ARC) through the Discovery Project grants DP190103388, DP200100704, DP210100990, and DE200101650 under the Discovery Early Career Researcher Award.

References

- Bearman, P.W. 1984, Vortex shedding from oscillating bluff bodies, *Annual Review of Fluid Mechanics*, **16**, 195–222.
- Bernitsas, M.M., Raghavan, K., Ben-Simon, Y. and Garcia, E.M.H. 2008, VIVACE (Vortex Induced Vibration Aquatic Clean Energy): A new concept in generation of clean and renewable energy from fluid flow, *Journal of Offshore Mechanics and Arctic Engineering*, **130**, 041101.
- Blevins, R.D. 1990, Flow-Induced Vibration, Van Nostrand Reinhold, Inc.
- Bonabeau, E., Dorigo, M. and Theraulaz, G. 1999, Swarm Intelligence: From Natural to Artificial Systems, Oxford University Press.
- Brunton, S.L., Noack, B.R. and Koumoutsakos, P. 2020, Machine Learning for Fluid Mechanics, *Annual Review of Fluid Mechanics*, **52**, 477-508.
- Lin, P., Hu, G., Li, C., Li, L., Xiao, Y., Tse, K.T. and Kwok, K.C.S. 2021, Machine learning-based prediction of crosswind vibrations of rectangular cylinders, *Journal of Wind Engineering and Industrial Aerodynamics*, **211**, 104549.
- Lo, J.C.C., Thompson, M.C., Hourigan, K. and Zhao, J. 2023, Effect of structural damping on flow-induced vibration of an elliptical cylinder, To be submitted to *Journal of Fluid Mechanics*.
- Lv, Y., Sun, L., Bernitsas, M.M and Sun, H. 2021, A comprehensive review of nonlinear oscillators in hydrokinetic energy harnessing using flow-induced vibrations, *Renewable and Sustainable Energy Reviews*, **150**, 111388.
- Ma, L., Resvanis T.L and Vandiver J.K. 2020, Using Machine Learning to Identify Important Parameters for Flow-Induced Vibration, in *International Conference on Offshore Mechanics and Arctic Engineering*, American Society of Mechanical Engineers, V004T04A011.
- Mei, Y.-F., Zheng, C., Aubry, N., Li, M.-G., Wu, W.-T. and Liu, X. 2021, Active control for enhancing vortex induced vibration of a circular cylinder based on deep reinforcement learning, *Physics of Fluids*, **33**, 103604.
- Naudascher, E. and Rockwell, D. 2005, Flow-Induced Vibrations: An Engineering Guide, Dover.
- Nemes, A, Zhao, J., Lo Jacono, D and Sheridan, J. 2012, The interaction between flow-induced vibration mechanisms of a square cylinder with varying angles of attack, *Journal of Fluid Mechanics*, **710**, 102–130.
- Shi, Y. and Eberhart, R. 1998, A modified particle swarm optimizer, in *1998 IEEE International Conference on Evolutionary Computation Proceedings. IEEE World Congress on Computational Intelligence (Cat. No. 98TH8360)*, IEEE, 69–73.
- Shi, Y. and Eberhart, R. 2001, Particle swarm optimization: developments, applications and resources, in *Proceedings of the 2001 Congress on Evolutionary Computation (IEEE Cat. No. 01TH8546)*, IEEE, 81–86.
- Slowik, A. and Kwasnicka, H. 2017, Nature inspired methods and their industry applications—Swarm intelligence algorithms, *IEEE Transactions on Industrial Informatics*, **14**, 1004–1015.
- Soti, A.K, Zhao, J., Thompson, M.C., Sheridan, J. and Bhardwaj, R. 2018, Damping effects on vortex-induced vibration of a circular cylinder and implications for power extraction, in *Journal of Fluids and Structures*, **81**, 289–308.
- Williamson, C.H.K. and Govardhan, R. 2004, Vortex-Induced Vibrations, *Annual Review of Fluid Mechanics*, **36**, 413–455.
- Zhao, J., Lo Jacono, D., Sheridan, J., Hourigan, K. and Thompson, M.C. 2018b, Experimental investigation of in-line flow-induced vibration of a rotating circular cylinder, *Journal of Fluid Mechanics*, **847**, 664–699.
- Zhao, J., Nemes, A., Lo Jacono, D., and Sheridan, J. 2018c, Branch/mode competition in the flow-induced vibration of a square cylinder, *Phil. Trans. R. Soc. Lond. A*, **376**, 20170243.
- Zhao, J., Leontini, J.S., Lo Jacono, D. and Sheridan, J. 2014, Fluid–structure interaction of a square cylinder at different angles of attack, *Journal of Fluid Mechanics*, **747**, 688–721.
- Zhao, J., Thompson, M.C. and Hourigan, K. 2023, Discovery of a hyper branch in flow-induced vibration of an elliptical cylinder, To be submitted to *Nature Communications*.