
Frequency-compensated PINNs for Fluid-dynamics Design Problems

Tongtao Zhang

Siemens Corporate Technology
Princeton, NJ 08540
tongtao.zhang@siemens.com

Biswadip Dey*

Siemens Corporate Technology
Princeton, NJ 08540
biswadip.dey@siemens.com

Pratik Kakkar

Siemens Corporate Technology
Princeton, NJ 08540
pratik.kakkar@siemens.com

Arindam Dasgupta

Siemens Corporate Technology
Princeton, NJ 08540
arindam.dasgupta@siemens.com

Amit Chakraborty

Siemens Corporate Technology
Princeton, NJ 08540
amit.chakraborty@siemens.com

Abstract

Incompressible fluid flow around a cylinder is one of the classical problems in fluid-dynamics with strong relevance with many real-world engineering problems, for example, design of offshore structures or design of a pin-fin heat exchanger. Thus learning a high-accuracy surrogate for this problem can demonstrate the efficacy of a novel machine learning approach. In this work, we propose a physics-informed neural network (PINN) architecture for learning the relationship between simulation output and the underlying geometry and boundary conditions. In addition to using a physics-based regularization term, the proposed approach also exploits the underlying physics to learn a set of Fourier features, i.e. frequency and phase offset parameters, and then use them for predicting flow velocity and pressure over the spatio-temporal domain. We demonstrate this approach by predicting simulation results over out of range time interval and for novel design conditions. Our results show that incorporation of Fourier features improves the generalization performance over both temporal domain and design space.

1 Introduction

The current approach for designing complex devices and systems, such as aero-dynamic surfaces and turbine components, typically involves an iterative interaction between design/operating space exploration and evaluation. However, high-fidelity fluid-dynamics simulations, which are necessary to evaluate the performance of design candidates under a variety of operating conditions, demand significant time and computational power. This limits the scope of the overall design optimization process and as a consequence, may lead to sub-optimal design choices. Applying machine learning algorithms to develop a fast and accurate surrogate for predicting simulation outcomes has the potential to significantly accelerate design evaluations thereby generating improved design choices.

*Corresponding author

In recent years, deep neural networks and representation learning [9] have improved significantly in accuracy and is widely-used in many application domains, such as image recognition [11], sequential decision making [26], and language comprehension [6]. These approaches, especially supervised learning, leverage datasets of inputs (e.g. pictures of people, historical data of a region’s average weather data, etc.) and outputs (e.g. identity of the individuals in a picture, and weather forecast, respectively) to learn their functional relationship using backpropagation. As inference in a neural network involves a single forward pass this provides an excellent opportunity to learn good-quality, low-cost surrogates for exploring the design spaces associated with fluid-dynamics problems.

To enable generalization beyond the training set, learning approaches incorporate appropriate inductive bias [5, 10] and promote representations which are *more plausible* in some sense. It typically manifests itself via a set of assumptions, which in turn can guide a learning algorithm to pick one hypothesis over another. The success in predicting an outcome for previously unseen data then depends on how well the inductive bias captures the ground reality. Inductive bias can be introduced as a prior in a Bayesian model, or via the choice of computational graphs and regularization terms in a neural network. In problems wherein laws of physics have a strong influence on the input-output relationship, generalization can be improved by leveraging underlying physics for defining the regularization term.

In this work, we incorporate a physics-based regularization term so that a learned surrogate conforms to the underlying physics governed by Navier-Stokes equations. We develop a PINN framework to infer how the velocity and pressure fields associated with the flow of an incompressible fluid around a cylinder depend on the underlying geometry and boundary condition - in particular, the size of the cylinder and inlet velocity, respectively. In addition, to capture the periodic nature of the solution which is governed by the Strouhal number [2], the proposed approach learns a set of Fourier features and a phase offset parameter as functions of the underlying geometry and boundary conditions. After the PINN has been trained, simulation results, i.e. velocity and pressure fields, for new design choices can be inferred in a fast, computationally inexpensive way. Thus, by enabling high-throughput evaluation of potential design candidates, this proposed approach provides a means to achieve better, more efficient design solutions. The key contributions of this work are as follows:

- We introduce a Fourier Feature Mapping (FFM) subnetwork within a PINN framework to yield better predictions about fluid flow around a cylinder. The FFM subnetwork learns frequency and phase offset parameters as a function of cylinder shape and inlet velocity.
- Subsequently, we use the learned surrogate to predict simulation outputs for novel design conditions and demonstrate the improvement in its generalization performance.

Related Work

ML-based Approaches for Fluid-dynamics Simulations: Use of machine learning algorithms in fluid-dynamics problems has drawn significant attention over the last few years. [12, 32] have shown that supervised learning, using large datasets of simulation results obtained from finite-element or finite-volume solvers, can build surrogates for predicting simulation results with high accuracy. [14, 20, 24, 31] have demonstrated that ML-based approaches can predict simulation results in mesh-free manner and incorporation of physics-based regularization in these formulations improves the quality of results by a significant margin. In addition, supervised learning has also been used to guide the discretization process in a data-driven way [3] or to learn efficient iterative solvers [13]. On the other hand, alternative approaches [7, 17, 19] based on self-supervised learning have also been proposed; they employ a neural network to approximate the solution of a differential equation and then use automatic differentiation to compute the loss function which is a quantitative measure on how well the dynamics (represented via a differential equation) and the initial/boundary conditions are enforced. As these approaches use the physics itself to generate training data, this line of work completely avoids the computationally expensive process for generating simulation datasets. It has also been shown that such neural network based approximations for a class of quasilinear, parabolic partial differential equations can converge to their true solutions with arbitrary accuracy [27].

Frequency Bias in Neural Networks: Frequency bias in neural networks is a relatively well-studied problem, with a body of work focusing on the relationship between frequency components present in a function and the speed at which neural networks learn them. [22] has demonstrated that neural networks with ReLU activation favors functions with low frequency components. [8, 18] have shown that deeper architectures are needed for a neural network to learn high-frequency functions. [4, 25] analyzed the learning dynamics in gradient descent and have shown that neural networks

learn low frequency functions much faster than high frequency functions. On the other hand, since the seminal work by Rahimi and Recht [23], multiple work have focused on accelerating learning algorithms by mapping the input to an appropriate feature space. A recent work [30] has shown that by using Fourier features as inputs to a multi-layer perceptron (MLP) can improve its capability to learn high frequency functions. Another work [28] has shown that use of periodic activation functions (e.g., \sin) can provide an efficient tool to learn representations in a large class of problems.

2 Frequency-compensated PINN for Fluid-dynamics Problems

2.1 Navier-Stokes Equations

In this work, we consider a classical fluid-dynamics problem, namely a cylinder in the cross flow [2]. In this flow configuration, an incompressible fluid passes around a cylinder. Then, by letting u , v denote the horizontal and vertical components of the velocity field and p denote the pressure field, the dynamics can be expressed as:

$$\left. \begin{aligned} \frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} + v \frac{\partial u}{\partial y} + \frac{\partial p}{\partial x} - \nu \left(\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} \right) &= 0 \\ \frac{\partial v}{\partial t} + u \frac{\partial v}{\partial x} + v \frac{\partial v}{\partial y} + \frac{\partial p}{\partial y} - \nu \left(\frac{\partial^2 v}{\partial x^2} + \frac{\partial^2 v}{\partial y^2} \right) &= 0 \\ \frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} &= 0 \end{aligned} \right\} \Psi(u, v, p) = 0, \quad (1)$$

where ν denotes the kinematic viscosity of the fluid under consideration and the functional Ψ describes the underlying partial differential equation. Moreover, by letting d_y and u_{inlet} denote the cylinder's diameter and inlet velocity, respectively, the Reynolds number for this flow can be expressed as $\text{Re} = (u_{inlet} * d_y) / \nu$. As we assume the kinematic viscosity to be 0.001, the *vortex shedding frequency* can be approximated as $0.21 * (1 - 21/\text{Re}) * (u_{inlet}/d_y)$ for the range of inlet velocity and cylinder diameter considered in this work [29].

2.2 PINN Architecture and the Fourier Component

Figure 1 illustrates the PINN architecture that we propose to learn a surrogate for the aforementioned problem. The input to the network are as follows: the spatial coordinates inside the domain (x, y) , the temporal dimension (t) , the inlet velocity (u_{inlet}) and the diameter of the cylinder (d_y) ; and the PINN predicts the velocity (u, v) and pressure (p) of the flow.




Figure 1: A diagram demonstrating the proposed framework. Arrows with dotted lines denote the data flow with regard to Fourier Feature Mapping.

As mentioned in the previous subsection, the velocity and pressure exhibit a periodic behavior along the temporal dimension. To capture this aspect, we use a set of Fourier features defined as

$$\text{Fourier Component} = [\sin(Ft + \phi), \cos(Ft + \phi)], \quad (2)$$

and employ an MLP to learn the frequency (F) and phase shift (ϕ) in the Fourier component as a function of u_{inlet} and d_y . These features are then fed into a second MLP subnetwork, along with the spatio-temporal coordinates (x, y, t) and design specifications (u_{inlet}, d_y) . The output from this

second subnetwork yield predictions on flow velocity and pressure. By letting \hat{u} , \hat{v} and \hat{p} denote the predicted velocity and pressure, respectively, we define the prediction error as:

$$L_{prediction} = \frac{1}{n} \sum_{i=1}^n [(u_i - \hat{u}_i)^2 + (v_i - \hat{v}_i)^2 + (p_i - \hat{p}_i)^2], \quad (3)$$

where n is the training dataset size. In addition, to enforce that the predicted values conform to the underlying physics governed by (1), we use the following regularization term

$$L_{PDE} = \frac{1}{n} \sum_{i=1}^n \|\Psi(u_i, v_i, p_i)\|^2. \quad (4)$$

Finally we define the following loss function for the PINN

$$L = L_{prediction} + \lambda L_{PDE}, \quad (5)$$

where the hyper-parameter λ maintains a balance between prediction accuracy and regularization.

3 Experiment

To evaluate performance of our proposed PINN framework, we train the network using simulation data corresponding to a handful of predefined geometry (d_y) and inlet velocity (u_{inlet}) combinations and use it to predict outputs for other combinations of geometry and inlet velocity values.

3.1 Dataset

We use FEniCS [1, 16], a finite-element solver, to create a dataset for the flow around a cylinder problem, as illustrated in Figure 2. We have a rectangular channel with length $L = 1.8$ and width $W = 0.4$, and an elliptical cylinder with fixed horizontal diameter $d_x = 0.1$ and vertical diameter d_y is placed inside the channel. The cylinder is placed at the vertical middle point and 0.2 from the inlet on the left. The region of interest is a rectangular region of length $X_{sample} = 1.5$ and width $Y_{sample} = 0.3$. We run FEniCS simulation for the time interval $[0s, 6s]$ (with a time resolution $0.05s$), and then sample the velocity and pressure data for this region of interest. In this work, we train and evaluate the PINN predict the output within this region of interest (shown in pink in Figure 2).




Figure 2: Geometry of the flow configuration. The pink area highlights the region of interest which is aligned in the vertical middle and is located in the right of the cylinder after an *offset* of 0.1 from its center.

In this work, we assume $u_{inlet} \in [0.8, 1.0]$ and $d_y \in [0.08, 0.11]$ and form the training set by running FEniCS simulation for the following 9 combinations (0.8, 0.08), (0.8, 0.09), (0.8, 0.10), (0.9, 0.08), (0.9, 0.10), (0.9, 0.11), (1.0, 0.08), (1.0, 0.09), and (1.0, 0.11) and then only taking the points from the time interval $[0s, 5s]$ with the interval of $0.1s$. Also, we keep the result for the following combinations of u_{inlet} and d_y as the validation set - (0.9, 0.09) and (0.9, 0.10); although the combination (0.9, 0.10) appears in the both these sets they do not overlap in the time direction. The rest of the points constitute the test set. With this, we have a dataset with training set of 517,350 instances, validation set of 206,700 instances and test set of 827,520 instances². If a geometry setting (i.e. a particular combination of u_{inlet} and d_y) appears in training set, we name it “Seen”, otherwise “Unseen”.

²For further details such as data point distribution among the geometry settings, please read the supplementary materials

3.2 Parameters and Settings

For the Fourier Feature Mapping (FFM) subnetwork, we use 3 fully-connected (FC) layers with 120 neurons on each layer and use Tanh activation function after the dropout layer. The final output size of the subnetwork is set as 10, and the first 5 outputs serve as frequency F in the Fourier component (2) while the rest 5 serve as phase shift ϕ .

For the second MLP subnetwork, we set up a network consisting of 10 FC layers with 120 neurons on each layer. The activation function for each FC layer is also chosen as Tanh.

In each step, besides the minibatches from the training set, we also randomly sample points from the rectangular region and we calculate the PDE loss (4) for those randomly sampled points so that the learned surrogate conforms to (1). The partial derivative equations are implemented using autograd toolbox from PyTorch [21]. The random points are drawn in the following domains: $x \in [0, 1.5]$, $y \in [0, 0.3]$, $t \in [0, 5]$, $u_{inlet} \in [0.8, 1]$ and $d_y \in [0.08, 0.11]$.

We use Adam Optimizer [15] to train the whole neural network with a learning rate of 0.001. The default run is 20,000 epoches with the minibatch size of 32768, and we adopt early-stopping strategy if the validation loss does not reduce.

3.3 Results

We show the errors and illustrate visualization output from our proposed framework given different geometry and component settings. Table 1 includes the MSE loss values at all the geometry and component settings which will be analyzed through the whole subsection.

	Seen/Covered	Unseen/Covered	Seen/Uncovered	Unseen/Uncovered
Full Component	1.09×10^{-4}	0.0269	0.0727	0.112
No-FFM	1.73×10^{-4}	0.347	0.0973	0.163
Strong-Reg	8.47×10^{-4}	0.0589	0.0453	0.0986
No-Reg/Overfit	2.51×10^{-3}	0.0431	0.0649	0.126

Table 1: Mean square error values of the testing sets and different experiment settings which will be discussed through Section 3.3. “Seen/Unseen” denotes whether the geometry settings appear in the training set (the data points in training and test set never overlap), “Covered/Uncovered” denotes whether the data is drawn from $t \in [0, 5]$ span.

3.3.1 Prediction within training geometry settings (Seen)

Figure 3a demonstrates the “worst” prediction output (the one possessing the highest MSE among the instances) of our proposed framework when using the geometry in the training set ($t = 4.65$, $u_{inlet} = 1.0$ and $d = 0.11$). It is clear that, within the training geometry settings, our proposed PINN framework provides accurate prediction on the inputs that do not appear at the training time stamps. The errors are barely noticeable on the visualization result and there is almost no phase shift.

3.3.2 Prediction outside training geometry settings (Unseen)

Figure 3b demonstrates the performance with the geometry settings that is unseen in the training set. We select the best (lowest MSE) prediction results and demonstrate it in Figure 3b. We can conclude that our framework generally works well in terms of recovering the data distribution in the sample space – or in terms of visualization, the shapes and artifacts in the images.

3.3.3 Ablation Settings, Challenges and Discussion

Beyond Time Span: In Table 1, we notice that points within time span $t \in [5, 6]$ (denoted as “Uncovered”) have lower performance or higher loss values. This is a major challenge in the proposed framework due to lack of further guidance in the training set and training process – no sampled points are drawn from time span $t \in [5, 6]$. However, as we illustrate the in Figure 4a, our proposed framework can still predict the flow velocity and pressure, albeit with some small phase shift.

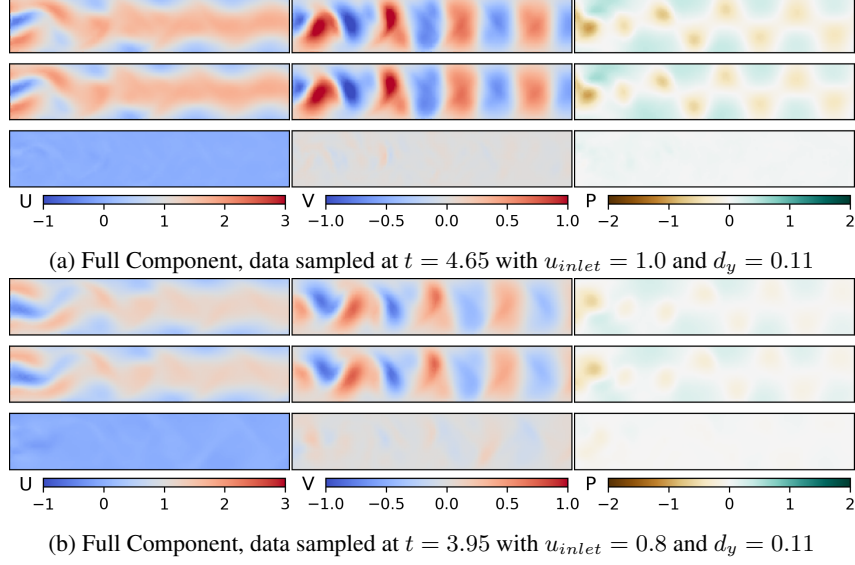


Figure 3: Visualization output as discussed in Section 3.3.1 and 3.3.2. In the same figure, the first row is the ground truth, the second row is the prediction and the third row is the error, other visualization figures have the same layout.

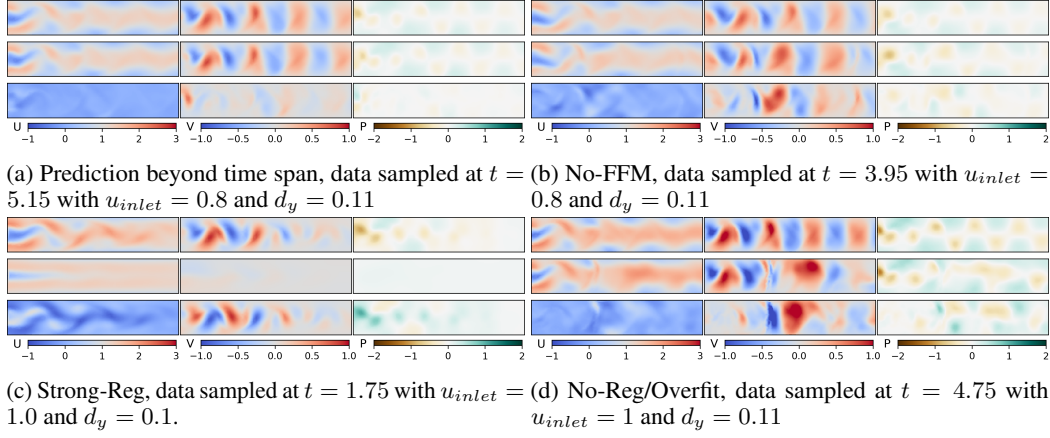


Figure 4: Visualization results which are discussed in Section 3.3.3

Fourier Feature Mapping: Figure 4b demonstrate the visualization of a framework that removes the FFM component – *i.e.*, x, y, t, u_{inlet} and d were directly fed into the second MLP. We use the same geometry setting and time stamp as the one in Figure 3b. From the output, we can conclude that, FFM provide crucial frequency and phase shift information in the output; and without it, the framework is not able to handle the complex frequencies and phase shifts among different combination of geometry settings.

Weighted Loss: In most of the experiments, the loss weight parameter λ was set as 0.001, and Figure 4c shows that a higher weighted PDE loss would inhibit the vorticity in the unseen settings. If we set $\lambda = 0$, meaning that we do not introduce PDE regularization in the framework, we encounter over fitting as demonstrated in Figure 4d.

Potential to Replace FeniCS-like Simulator: When we conduct FeniCS simulation on one geometry setting, with a time step of 0.0001 on a machine equipped with two Intel Xeon CPU E5-2620 v4 CPUs and 4 Nvidia Titan Xp, we spend more than 24,000 seconds in simulating data over the time interval $[0s, 6s]$. With our proposed framework, we only need 50 time steps as anchor points from

FeniCS, and within 15,000 seconds (even we do not use early-stopping) we are able to acquire the same amount of data with comparable quality.

4 Conclusion

In this paper, we proposed a frequency-compensated PINN framework which can build high-accuracy surrogate for predicting simulation result in fluid dynamics design problems. In particular, we introduced and leveraged a Fourier feature mapping subnetwork to capture the periodicity present in the flow velocity and pressure; to conform with the underlying physics governed by the Strouhal number, this subnetwork learns the Fourier components as a functions of inlet velocity and cylinder size. Our results show that these Fourier features improve generalization in spatial and temporal domain as well for novel geometry settings. Future work would explore how this framework can be further extended to address fluid dynamics problems with more complex geometry and boundary conditions.

References

- [1] M. S. Alnæs, J. Blechta, J. Hake, A. Johansson, B. Kehlet, A. Logg, C. Richardson, J. Ring, M. E. Rognes, and G. N. Wells. The fenics project version 1.5. *Archive of Numerical Software*, 3(100), 2015. doi: 10.11588/ans.2015.100.20553.
- [2] J. D. Anderson and J. Wendt. *Computational fluid dynamics*, volume 206. Springer, 1995.
- [3] Y. Bar-Sinai, S. Hoyer, J. Hickey, and M. P. Brenner. Learning data-driven discretizations for partial differential equations. *Proceedings of the National Academy of Sciences*, 116(31):15344–15349, 2019.
- [4] R. Basri, M. Galun, A. Geifman, D. Jacobs, Y. Kasten, and S. Kritchman. Frequency bias in neural networks for input of non-uniform density, 2020.
- [5] J. Baxter. A model of inductive bias learning. *Journal of Artificial Intelligence Research*, 12:149–198, 2000.
- [6] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [7] V. Dwivedi, N. Parashar, and B. Srinivasan. Distributed physics informed neural network for data-efficient solution to partial differential equations. *arXiv preprint arXiv:1907.08967*, 2019.
- [8] R. Eldan and O. Shamir. The power of depth for feedforward neural networks. In *Conference on learning theory*, pages 907–940, 2016.
- [9] I. Goodfellow, Y. Bengio, A. Courville, and Y. Bengio. *Deep learning*, volume 1. MIT press Cambridge, 2016.
- [10] D. Haussler. Quantifying inductive bias: AI learning algorithms and Valiant’s learning framework. *Artificial Intelligence*, 36(2):177–221, 1988.
- [11] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [12] O. Hennigh. Lat-net: compressing lattice boltzmann flow simulations using deep neural networks. *arXiv preprint arXiv:1705.09036*, 2017.
- [13] J.-T. Hsieh, S. Zhao, S. Eismann, L. Mirabella, and S. Ermon. Learning neural PDE solvers with convergence guarantees. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=rklaWn0qK7>.
- [14] C. M. Jiang, S. Esmaeilzadeh, K. Azizzadenesheli, K. Kashinath, M. Mustafa, H. A. Tchelepi, P. Marcus, A. Anandkumar, et al. Meshfreeflownet: A physics-constrained deep continuous space-time super-resolution framework. *arXiv preprint arXiv:2005.01463*, 2020.
- [15] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [16] A. Logg, K.-A. Mardal, G. N. Wells, et al. *Automated Solution of Differential Equations by the Finite Element Method*. Springer, 2012. ISBN 978-3-642-23098-1. doi: 10.1007/978-3-642-23099-8.

- [17] L. Lu, X. Meng, Z. Mao, and G. E. Karniadakis. Deepxde: A deep learning library for solving differential equations. *arXiv preprint arXiv:1907.04502*, 2019.
- [18] G. F. Montufar, R. Pascanu, K. Cho, and Y. Bengio. On the number of linear regions of deep neural networks. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2924–2932. 2014.
- [19] M. A. Nabian and H. Meidani. A deep learning solution approach for high-dimensional random differential equations. *Probabilistic Engineering Mechanics*, 57:14–25, 2019.
- [20] M. A. Nabian and H. Meidani. Physics-driven regularization of deep neural networks for enhanced engineering design and analysis. *Journal of Computing and Information Science in Engineering*, 20(1), 2020.
- [21] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- [22] N. Rahaman, A. Baratin, D. Arpit, F. Draxler, M. Lin, F. Hamprecht, Y. Bengio, and A. Courville. On the spectral bias of neural networks. volume 97 of *Proceedings of Machine Learning Research*, pages 5301–5310, Long Beach, California, USA, 09–15 Jun 2019. PMLR.
- [23] A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *Advances in neural information processing systems*, pages 1177–1184, 2008.
- [24] M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- [25] B. Ronen, D. Jacobs, Y. Kasten, and S. Kritchman. The convergence rate of neural networks for learned functions of different frequencies. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 4761–4771. 2019.
- [26] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, et al. Mastering the game of Go without human knowledge. *Nature*, 550(7676):354–359, 2017.
- [27] J. Sirignano and K. Spiliopoulos. Dgm: A deep learning algorithm for solving partial differential equations. *Journal of computational physics*, 375:1339–1364, 2018.
- [28] V. Sitzmann, J. N. Martel, A. W. Bergman, D. B. Lindell, and G. Wetzstein. Implicit neural representations with periodic activation functions. *arXiv preprint arXiv:2006.09661*, 2020.
- [29] B. M. Sumer et al. *Hydrodynamics around cylindrical structures*, volume 26. World scientific, 2006.
- [30] M. Tancik, P. P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. T. Barron, and R. Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *NeurIPS*, 2020.
- [31] R. Wang, K. Kashinath, M. Mustafa, A. Albert, and R. Yu. Towards physics-informed deep learning for turbulent flow prediction. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1457–1466, 2020.
- [32] Q. Wei, I. Akrotirianakis, A. Dasgupta, and A. Chakraborty. Learn to learn: Application to topology optimization. *Smart and Sustainable Manufacturing Systems*, 2(1):250–260, 2018.

A Data Distribution

Table 2 demonstrates the distribution of the training, validation and test set. If a geometry setting appears in training set, we name it “Seen”, otherwise “Unseen”.

$\{u_{inlet}, d_y\}$	Train	Validation	Test
$\{0.8, 0.08\}$	57,700	/	69,240
$\{0.8, 0.09\}$	57,500	/	69,000
$\{0.8, 0.10\}$	57,400	/	68,880
$\{0.8, 0.11\}$	/	/	137,760
$\{0.9, 0.08\}$	57,700	/	69,240
$\{0.9, 0.09\}$	/	138,000	/
$\{0.9, 0.10\}$	57,250	68,700	/
$\{0.9, 0.11\}$	57,400	/	68,880
$\{1.0, 0.08\}$	57,500	/	69,240
$\{1.0, 0.09\}$	57,400	/	69,000
$\{1.0, 0.10\}$	/	/	137,400
$\{1.0, 0.11\}$	57,500	/	68,880
Sum	517,350	206,700	827,520

Table 2: Statistics and split of training, validation and test sets. Due to internal mechanics of FeniCS, such as mesh grid settings, there are slight differences in the numbers of data points among different geometry settings.