



HAL
open science

Understanding the training of PINNs for unsteady flow past a plunging foil through the lens of input subdomain level loss function gradients

Rahul Sundar, Didier Lucor, Sunetra Sarkar

► To cite this version:

Rahul Sundar, Didier Lucor, Sunetra Sarkar. Understanding the training of PINNs for unsteady flow past a plunging foil through the lens of input subdomain level loss function gradients. 10th International and 50th National Conference on Fluid Mechanics and Fluid Power (FMFP), Dec 2023, Jodhpur, France. hal-04463178

HAL Id: hal-04463178

<https://hal.science/hal-04463178v1>

Submitted on 7 Mar 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Understanding the training of PINNs for unsteady flow past a plunging foil through the lens of input subdomain level loss function gradients

Rahul Sundar^{1*}, Didier Lucor², and Sunetra Sarkar¹

^{1*}Department of Aerospace Engineering, Indian Institute of Technology Madras, Chennai 600036, India

²Laboratoire Interdisciplinaire des Sciences du Numérique LISN-CNRS, Université Paris-Saclay, Orsay 91403, France

*rahulsundar@smail.iitm.ac.in; didier.lucor@lisn.upsaclay.fr; and sunetra.sarkar@gmail.com

ABSTRACT Recently immersed boundary method inspired physics-informed neural networks (PINNs) including the moving boundary-enabled PINNs (MB-PINNs) have shown the ability to accurately reconstruct velocity and recover pressure as a hidden variable for unsteady flow past moving bodies. Considering flow past a plunging foil, MB-PINNs were trained with global physics loss relaxation and also in conjunction with a physics-based undersampling method, obtaining good accuracy. The purpose of this study was to investigate which input spatial subdomain contributes to the training under the effect of physics loss relaxation and physics-based undersampling. In the context of MB-PINNs training, three spatial zones: the moving body, wake, and outer zones were defined. To quantify which spatial zone drives the training, two novel metrics are computed from the zonal loss component gradient statistics and the proportion of sample points in each zone. Results confirm that the learning indeed depends on the combined effect of the zonal loss component gradients and the proportion of points in each zone. Moreover, the dominant input zones are also the ones that have the strongest solution gradients in some sense.

Keywords: Immersed boundaries, Physics informed neural networks, surrogate modeling, unsteady flows, plunging foil

I. INTRODUCTION

Physics-informed neural networks have become promising for their use in solving complex inverse problems such as hidden physics recovery, data-driven equations discovery, uncertainty quantification, Plain vanilla PINNs ([1] are however difficult to train when the systems exhibit strong spatiotemporal gradients. Hence, recently, authors have proposed different strategies to train PINNs better, such as adaptive sampling, modified architectures, static and dynamic loss weighting.

Surrogate modeling of unsteady flow past moving boundaries such as flapping wings become challenging if the underlying high-fidelity simulation data used for such an endeavour is obtained using the immersed boundary method [2], [3]. This is because, IBM uses a fixed Eulerian background grid for the fluid, whereas the solid boundary is described by a set of Lagrangian markers. At any time

instant, there exist eulerian grid cells bounded by the solid boundary which consists of fictitious flow field data.

Flows past flapping wings are often characterized by strong flow-field gradients which makes it challenging for PINNs to train. Hence, in a recent work, a moving boundary enabled PINNs (MB-PINNs) [4] inspired by the IBM was proposed, where a fixed Eulerian grid for the fluid domain was considered discarding the solid region points a any given time and the no-slip boundary condition was enforced directly on the Lagrangian markers unlike in discrete forcing IBM [5]. A combination of global physics loss relaxation and vorticity cutoff based under sampling was shown to improve data efficiency while maintaining good accuracy in velocity reconstruction and simultaneous pressure recovery as a hidden variable.

Since loss component weighting improves the loss balancing globally as reported in [6], it would be worthwhile to investigate if there exist local imbalances in the spatial domains of interest. This would also in a way determine which spatial region drives the training. One way of understanding how PINNs train is to look into the layer-wise loss component gradients as in [6] to understand the effect of competing objectives on the training. While these gradients are often visualized for an entire full batch setting, it is not visualized however in a localized context to understand the contributions from respective spatiotemporal zones. Visualizing the layer-wise gradients obtained for the loss components over specific spatiotemporal zones would highlight any imbalances in the gradient updates and also validate the need for localized weighting strategies further like in the self-adaptive weighting or residual-based attention approaches [7], [8].

Hence, the current study's objectives are to investigate how loss contributions spatially are affected by physics loss relaxation combined with a physics-based undersampling and devise metrics to quantify which spatial zone drives the training of the network. As an example, the flow past a plunging airfoil at a low Reynolds number is considered following [4]. From [4], three MB-PINN cases are considered with one being a baseline without any loss weighting or under-sampling and the other two with loss relaxation and undersampling.

The outline of the paper is as follows. The methodology of MB-PINNs and understanding the loss component gradi-

ents is presented in section II. The numerical experiments and results are discussed in section III and finally the conclusions are presented in section IV.

II. METHODOLOGY

In this section, the problem setup, MB-PINN formulation, zonal splitting of loss component gradients, and metrics to quantify the zonal imbalances and determine the spatial zone driving the training will be discussed.

A. The plunging foil system

In the present study, incompressible unsteady flow past a harmonically plunging elliptic foil is considered as an example. The flow around the foil is governed by the two-dimensional incompressible Navier-Stokes equations

$$\partial_t \hat{\mathbf{u}} + \hat{\mathbf{u}} \cdot \nabla (\hat{\mathbf{u}}) = -\nabla \hat{p} + \frac{1}{Re} \nabla^2 \hat{\mathbf{u}}, \quad \text{and} \quad (1)$$

$$\nabla \cdot \hat{\mathbf{u}} = 0, \quad (2)$$

where Re is the Reynolds number, and $\hat{\mathbf{u}}, \hat{p}$ are velocity and pressure fields satisfying the boundary conditions as shown in the problem setup schematic in Fig. 1, respectively. On the moving boundary Γ_{IB} especially, a no-slip boundary condition dependent on the plunging kinematics is to be satisfied. Here, the plunging kinematics is modeled as follows

$$y(t) = h_0 \cos(\omega_h t), \quad \text{and} \quad (3)$$

$$\dot{y}(t) = -\omega_h h_0 \sin(\omega_h t), \quad (4)$$

where t, h_0 and ω_h are dimensional time, the plunging amplitude and frequency, respectively. The non-dimensional amplitude $h = h_0/c$ and reduced frequency $k = \omega_h c/U_\infty$ are defined aligning with literature [9], [10], where, $U_\infty = 1$ is the free stream velocity and $c = 1$ is the chord length of the foil.

In the present study, the example from [4] with $Re = 500, k = 2\pi$ and $h = 0.16$ is chosen aligning with [10] for gradient-based analysis. The training data has been generated using a discrete forcing IBM-based unsteady flow solver [5] where, additional momentum forcing (f) and source/sink (q) terms are added to Eqs. (1) and (2), respectively to satisfy the no-slip boundary condition.

B. Moving boundary enabled PINNs (MB-PINNs)

Inspired by the immersed boundary method, an immersed boundary aware framework using PINNs for surrogate modeling of unsteady flows past moving boundaries was explored by Sundar *et al.* [4]. A truncated spatial domain excluding solid region points Ω_f^r was chosen as shown in Fig. 1 to train the moving boundary enabled PINN(MB-PINN) models in [4]. It was shown that the MB-PINN (see Fig. 2) was efficacious against solving velocity reconstruction and simultaneous pressure recovery given IBM data. A feed forward neural network with Swish activation as a backbone was used in [4] to map the spatio-temporal coordinates to the corresponding velocity and pressure predictions. The loss function \mathcal{L} can be written as follows

$$\mathcal{L} = \lambda_{Bulk} \mathcal{L}_{Bulk} + \lambda_{BC} \mathcal{L}_{BC} + \lambda_{IC} \mathcal{L}_{IC} + \lambda_{IB} \mathcal{L}_{IB} + \lambda_{Phy} \mathcal{L}_{Phy}. \quad (5)$$

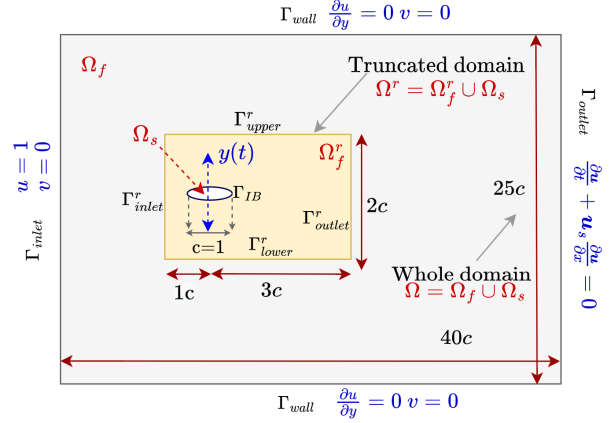


Figure 1: Problem setup (Image adapted from Sundar *et al.*) [4].

Here, $\mathcal{L}_\#$ with the subscripts $\# = \{\text{Bulk, BC, IC, IB, Phy}\}$ denote the loss contributions from mean squared errors in interior bulk velocity data, boundary condition, initial conditions, no-slip boundary condition on the immersed boundary and the physics constraints, respectively and $\lambda_\#$ denotes the corresponding loss weights. Here, the boundary conditions include Dirichlet values at the inlet, upper and lower boundaries of the truncated computational domain Ω_f^r (see Fig. 1).

Throughout the study, $\lambda_{Bulk} = \lambda_{BC} = \lambda_{IC} = \lambda_{IB} = 1$, and the physics loss alone is relaxed. In addition to physics loss relaxation, a vorticity cutoff based sampling was proposed in [4]. Here, a vorticity cutoff $|\omega_z^*|$, is chosen and all those fluid data points retained such that $|\omega_z| \geq |\omega_z^*|$. Then based on a percentage sampling ratio $S_{\omega_z} = \frac{N_{\omega_z \geq |\omega_z^*|}^{sample}}{N_{|\omega_z| < |\omega_z^*|}} \times 100$, the remaining fluid points are undersampled.

Considering K_{MB} mini batches of the data, the network weights and biases together represented by θ are updated using the back propagated loss component gradients as follows

$$\theta^{i+1} = \theta^i - \eta_i \frac{1}{N_{MB}} \sum_{i'=iN_{MB}+1}^{(i+1)N_{MB}} \left(\lambda_{Bulk} \nabla_{\theta} \mathcal{L}_{Bulk}^{i'} + \lambda_{BC} \nabla_{\theta} \mathcal{L}_{BC}^{i'} + \lambda_{IB} \nabla_{\theta} \mathcal{L}_{IB}^{i'} + \lambda_{IC} \nabla_{\theta} \mathcal{L}_{IC}^{i'} + \lambda_{Phy} \nabla_{\theta} \mathcal{L}_{Phy}^{i'} \right), \quad (6)$$

where η_i is the learning rate corresponding to i^{th} iteration and each epoch corresponds to $K_{MB} = N_{train}/N_{MB}$ iterations. Here, the MB-PINN models were trained using the stochastic gradient descent based optimisation algorithm ADAM [11]. Unless otherwise specified, a mini-batch size of $N_{MB} = 1500$ was chosen with a learning rate step decay $lr = [1e - 03, 53 - 04, 1e - 04]$ where each learning rate step consisted of $5e05$ training iterations.

The loss components $\mathcal{L}_{IB}, \mathcal{L}_{BC}$ don't depend on the

interior spatial coordinates, and \mathcal{L}_{IC} considers a data snapshot only at $t/T = 0$. They play an important role in ensuring the hidden variables are recovered appropriately [12], [13]. However, \mathcal{L}_{Bulk} and \mathcal{L}_{Phy} are computed from spatial coordinates interior to Ω_f^r . Given that spatially, there are zones where strong, weak or no vortices exist, it would thus be interesting to investigate the contributions from different spatial regions internally in the domain.

C. Zonal splitting of layer-wise gradients

While global loss weighting improves training and achieves global loss component balancing, that might not be the case locally. This could be due to varied contributions from each spatial region leading to different effective learning rates in each spatial region. To analyse this, once the MB-PINNs are trained or even at any intermediate learning stage, the layer-wise gradients can be computed overall for samples over the entire domain and for points sampled from select spatial zones. In the present study, three spatial zones are defined (see Fig. 3), namely, the moving body zone ($Z1 \subset \Omega_f^r$ such that $-1 < x < 1$ and $-0.5 < y < 0.5$), the wake zone ($Z2 \subset \Omega_f^r$ such that $1 < x$ and $-0.5 < y < 0.5$), and the outer zone ($Z3 \subset \Omega_f^r$ such that $Z3 = \Omega_f^r - Z1 \cup Z2$). Note that the overall spatial domain is Ω_f^r which is $\Omega^r = [-1c, 3.5c] \times [-1c, 1c]$ minus the solid region points at any time $t \in [0, 2]$. Also note that, zonal splitting is not carried out in temporal domain.

Given training velocity data $\hat{\mathbf{u}}$ at these spatial coordinates across all time, the components \mathcal{L}_{Bulk} and \mathcal{L}_{Phy} in Eq. 5 can hence be split into contributions from respective spatial-zones as follows

$$\mathcal{L}_{Bulk} = \frac{1}{N_{MB}} \left(\sum_{i=1}^3 \left(\sum_{j=1}^{N_{Zi}} \|\mathbf{u}^j - \hat{\mathbf{u}}^j\|^2 \right) \right) \quad (7)$$

$$= \left(\sum_{i=1}^3 \left(\frac{N_{Zi}}{N_{MB}} \sum_{j=1}^{N_{Zi}} \frac{\|\mathbf{u}^j - \hat{\mathbf{u}}^j\|^2}{N_{Zi}} \right) \right) \quad (8)$$

$$= \left(\sum_{i=1}^3 \left(p_{Zi} \sum_{j=1}^{N_{Zi}} \frac{\|\mathbf{u}^j - \hat{\mathbf{u}}^j\|^2}{N_{Zi}} \right) \right) \quad (9)$$

$$= \left(\sum_{i=1}^3 \left(p_{Zi} \mathcal{L}_{Bulk}^{Zi} \right) \right) \quad (10)$$

It is seen that \mathcal{L}_{Bulk} is a weighted sum of zonal mean squared errors where p_{Zi} for $i = 1, 2, 3$ is the proportion of sample points from zone Zi with respect to the overall mini-batch sample. Similarly, \mathcal{L}_{Phy} can be split into respective contributions from the spatial zones as follows

$$\mathcal{L}_{Phy} = \left(\sum_{i=1}^3 \left(p_{Zi} \sum_{j=1}^{N_{Zi}} \frac{\|\mathbf{r}(\mathbf{u}^j)\|}{N_{Zi}} \right) \right) \quad (11)$$

The above expressions can be simplified as

$$\mathcal{L}_* = \left(\sum_{i=1}^3 \left(p_{Zi} \mathcal{L}_*^{Zi} \right) \right) \quad (12)$$

for $* = \{\text{Bulk, Phy}\}$. Now the gradient vector $\nabla_{\theta} \mathcal{L}_*$ required to be computed at every training iteration to update the network parameters can be further expressed in terms of the zonal contributions as

$$\nabla_{\theta} \mathcal{L}_* = \left(\sum_{i=1}^3 \left(p_{Zi} \nabla_{\theta} \mathcal{L}_*^{Zi} \right) \right). \quad (13)$$

The above zonal splitting further indicates that in addition to loss component weighting, the data and collocation points sampling can modify the contributions from spatial zones or even spatio-temporal zones of interest. Hence, the zonal contributions can't be investigated independent of the sample proportions as the overall loss gradient is the weighted sum of the gradients obtained over each zone.

D. Metrics to diagnose zonal imbalances

Visualising and comparing the distribution of layerwise loss component gradients often sheds light on the extent of balance between the components [6] globally or locally. In addition, as seen previously, since gradient updates also depend on the proportion of zone specific coordinates with respect to the overall sample $p_{Z\#}$, this has to be weighed in while quantifying the relative contributions from each spatial zone. To quantify which zone drives the training for components \mathcal{L}_* with $* = \{\text{Bulk, Phy}\}$, the following zonal proportion weighted relative gradient statistics based metrics are thus computed

$$\mu_*^{Zi} = p_{Zi} \frac{|\nabla_{\theta} \mathcal{L}_*^{Zi}|}{|\nabla_{\theta} \mathcal{L}_*|} \quad (14)$$

$$\sigma_*^{Zi} = p_{Zi} \frac{\text{std}\{|\nabla_{\theta} \mathcal{L}_*^{Zi}|\}}{\text{std}\{|\nabla_{\theta} \mathcal{L}_*|\}}. \quad (15)$$

Once MB-PINNs are trained, the mean and standard deviation of the gradient magnitudes $|\nabla_{\theta} \mathcal{L}_*|$, $\text{std}\{|\nabla_{\theta} \mathcal{L}_*|\}$ are computed over the network parameters for a mini-batch sample from the entire spatio-temporal domain and also for samples from respective spatial zones. The metrics μ_*^{Zi} , σ_*^{Zi} for $i = \{Z1, Z2, Z3\}$, and $* = \{\text{Bulk, Phy}\}$ are then computed.

III. RESULTS AND DISCUSSION

In the present study, three MB-PINN test cases from [4] are considered for analysis; a model without any physics loss relaxation or under sampling (case 1), a model with physics loss relaxation alone (case 2), and a model with both physics loss relaxation and vorticity cutoff based undersampling (case 3). The true and predicted velocity x-component, associated point-wise errors, and vorticity contours for these cases are presented in Fig. 4 where the errors are prominent in zone $Z1$ and $Z2$ for case 1 but they almost vanish for cases 2 and 3 (see Figs. 4(c-e)). The accuracy details are also presented in Table 3. To further understand which spatial zones in the flow contribute to the training and determine if there exist spatial imbalances, the loss component wise gradient distributions are analysed. The metrics μ_*^{Zi} and σ_*^{Zi} for $i = 1, 2, 3$ and $* = \{\text{Bulk, Phy}\}$ described in Eqs. (14) and (15) are presented in Table 4 for all the test cases discussed below.

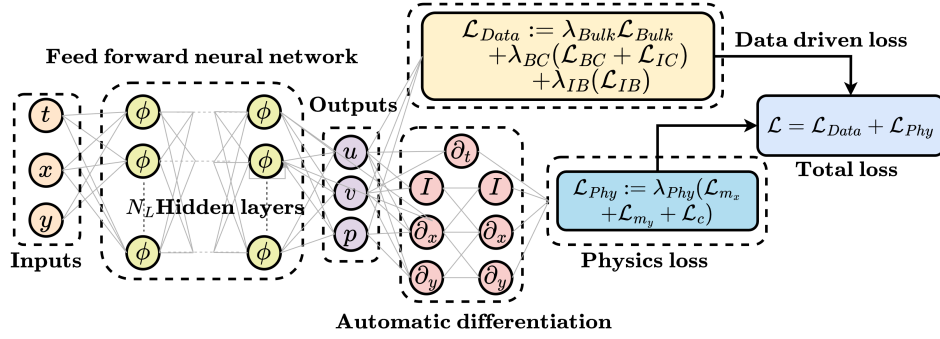


Figure 2: MB-PINN schematic




Figure 3: Spatial grids of (a) CI and (b) CI-S5 data sets and the corresponding (c-d) zonal splitting of the grids into Z1- moving body, Z2 - wake and Z3 - outer zones, respectively

Table 1: Training and testing data set resolution within the truncated domain considered in this study. Ref-IBM and Ref-ALE are the high resolution testing data-sets generated using a IBM and an ALE solver, respectively. Further details about the 'Ref-*' data-sets can be found in [4].

Data sets	N_x	N_y	N_t	$\Delta t/T$	$(N_x \times N_y \times N_t)$
CI	270	120	41	0.05	1.3284e06
CI-S5	-	-	-	0.05	2.648e05
Ref-IBM	651	500	81	0.025	2.6365e07
Ref-ALE	-	-	81	0.025	4.05e06

A. Case 1: Baseline ($\lambda_{Phys} = 1$)

It is observed that the first and last layer gradient distributions overall (see Figs. 5(a) and 5(d)) indicate an imbalance in contributions more prominent in the first layer. The $\nabla_{\theta} L_{Phys}$ distributions across all zones are at least one order higher than that of $\nabla_{\theta} L_{Bulk}$ (see Figs. 6(a) and 6(d))

Table 2: Training data-sets and associated proportion p_{Zi} of data points from each spatial zone Zi with $i = 1, 2, \text{ and } 3$ for a mini-batch sample at $1e06$ iterations.

Data set	N_{Bulk}	p_{Z1}	p_{Z2}	p_{Z3}
CI	1.2915e06	2.175e-01	2.857e-01	4.967e-01
CI-S5	2.6481e05	2.936e-01	5.841e-01	1.223e-01

Table 3: Accuracy of MB-PINN for different relaxation coefficients evaluated on Ref-IBM and Ref-ALE datasets. Best performing models are highlighted in bold-faced fonts.

Test case	λ_{fluid}	S_{ω_z} (in %)	aMAE	arRMSE (in %)
Case 1	1	100	1.37e-01	22.96
Case 2	0.001	100	2.07e-02	3.22
Case 3	0.01	5	2.25e-02	3.22

and Figs. 7(a) and 7(d), respectively). Although the gradients have a higher μ and σ (subscripts dropped for simplicity) in Z1 (see Table 4) the gradients for the overall domain Ω_f still receive significant contributions from Z2 and Z3 (see Fig. 6(a) and 6(d)) since $p_{Z1} < p_{Z2} < p_{Z3}$.

B. Case 2: With physics loss relaxation ($\lambda_{Phys} = 0.001$)
 With global physics loss relaxation, the MB-PINN model trained on the CI dataset was reported to be optimal for $\lambda_{Phys} = 0.001$. It is seen that the gradient magnitudes are much lower than that of the baseline case indicating some sort of vanishing gradient problem (see Figs. 5(b) and 5(e)). Even though the gradient magnitudes are very low, the loss component gradients are relatively well balanced across all the zones unlike in case 1 (see Figs. 6(b) and 6(e), and Figs. 7(b) and 7(e), respectively). Here, as seen in the Table 4, $\nabla_{\theta} L_{Bulk}$ is dominated by contributions from Z2, while $\nabla_{\theta} L_{Phys}$ receives most of its contributions from Z1. In spite of the good accuracy, it is seen that the gradients almost vanish in the last layer as opposed to the first layer.

C. Case 3: With global physics loss relaxation and vorticity cutoff based under sampling ($\lambda_{Phys} = 0.01, S_{\omega_z} = 5\%$)

Combining the physics loss relaxation with a physics based vorticity cutoff sampling was identified as a useful




Figure 4: Comparison of (a-b) true and (c-e) MB-PINN predicted velocity x -component, corresponding point-wise maximum value normalised absolute errors and vorticity contours at $t/T = 1.0$.

strategy in [4] to reduce the data requirement while still obtaining similar accuracy as in case 2 (see Table 3). It is observed that in case 3 the magnitudes of $\nabla_{\theta}\mathcal{L}_{Bulk}$ and $\nabla_{\theta}\mathcal{L}_{Phy}$ are higher and they don't vanish unlike in case 2 (see Figs. 5(c) and 5(f)). In a way, since vorticity cutoff based under sampling retains more data points only in those regions with strong gradients, it is possible that this prevents the gradients from vanishing. Unlike Case 1 and Case 2, due to selective under sampling, $p_{Z3} < p_{Z1} < p_{Z2}$ which in turn affects the gradient back propagation. As a result, in Table 4 it is seen that $\nabla_{\theta}\mathcal{L}_{Bulk}$ and $\nabla_{\theta}\mathcal{L}_{Phy}$ are both dominated by contributions from $Z1$ followed by that from $Z2$. The first and last layer gradient distributions as well indicate that across all zones, the gradients are balanced with $Z1$ driving the training as seen by the gradient distributions on Ω_f following $Z1$.

IV. CONCLUSIONS

In this paper, to understand the training of MB-PINNs for flow past a plunging foil, a novel zonal splitting methodology and gradient statistics-based metrics were proposed to analyze spatial imbalances in loss component contributions. Three zones consisting of the moving body, the wake and the outer far-field were considered. Three test cases from [4] were considered for analysis with and without physics loss relaxation and vorticity cutoff based under-sampling. The analysis confirmed the existence of zonal imbalances and also determined which zone dictated training. In the test case with the relaxation of the physics loss alone, the

Table 4: Zonal relative gradient statistics across test cases. Dominant zones are highlighted in bold.

Zone	$Z\#$ μ_{Bulk}	$Z\#$ σ_{Bulk}	$Z\#$ μ_{Phy}	$Z\#$ σ_{Phy}
Case 1: $\lambda_{Phy} = 1$				
Z1	7.932e-01	7.466e-01	1.256	2.153
Z2	2.179e-01	2.669e-01	1.292e-01	1.368e-01
Z3	2.292e-01	2.459e-01	6.503e-02	5.801e-02
Case 2: $\lambda_{Phy} = 0.001$				
Z1	3.207e-01	4.769e-01	5.565e-01	5.608e-01
Z2	7.552e-01	8.477e-01	6.906e-02	5.608e-02
Z3	1.297e-01	1.676e-01	9.610e-3	5.823e-03
Case 3: $\lambda_{Phy} = 0.01, S_{\omega_z} = 5\%$				
Z1	7.054e-01	1.194	4.929	7.555
Z2	4.872e-01	4.392e-01	1.359e-01	1.073e-01
Z3	7.959e-03	5.917e-03	1.432e-03	7.506e-04

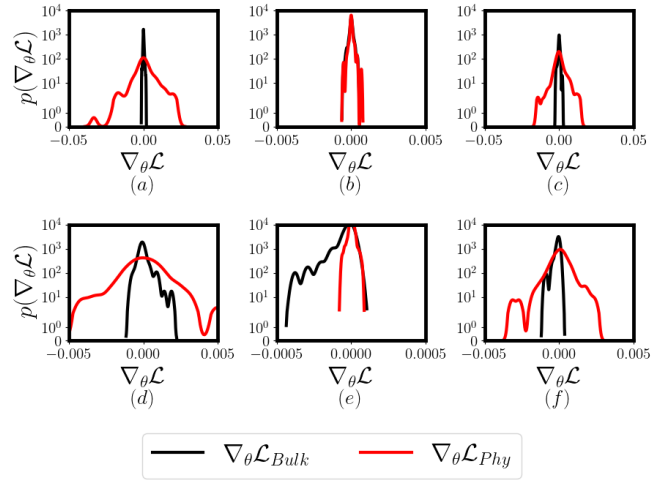


Figure 5: First layer (top row) and last layer (bottom row) $\nabla_{\theta}\mathcal{L}_{Bulk}$ and $\nabla_{\theta}\mathcal{L}_{Phy}$ distributions obtained after $1e06$ training iterations for (a,d) Case 1: $\lambda_{Phy} = 1$, (b,e) Case 2: $\lambda_{Phy} = 0.001$ and $S_{\omega_z} = 100\%$ and (c,f) Case 3: $\lambda_{Phy} = 0.01$ and $S_{\omega_z} = 5\%$.

wake zone dominated the gradient updates coming from the loss of data. Whereas, the moving body zone dominated the updates from physics loss. However, for the test case with both physics loss relaxation and physics-based under-sampling, the moving body zone dominated the gradient updates for both the loss components. It was also observed that a vorticity cutoff based under-sampling alleviates the problem of vanishing gradients as opposed to the case of no under-sampling, further validating the importance of selective physics-based sampling. It is envisioned that quantifying the imbalanced contribution from spatial or even more generally input subdomains would enable the design of effective loss component weighting strategies. This approach can be extended effectively for other problems as well to evaluate imbalances in contributions from the input subdomains of interest. Moreover, being agnostic to the model architecture, this approach can be applied to physics-based

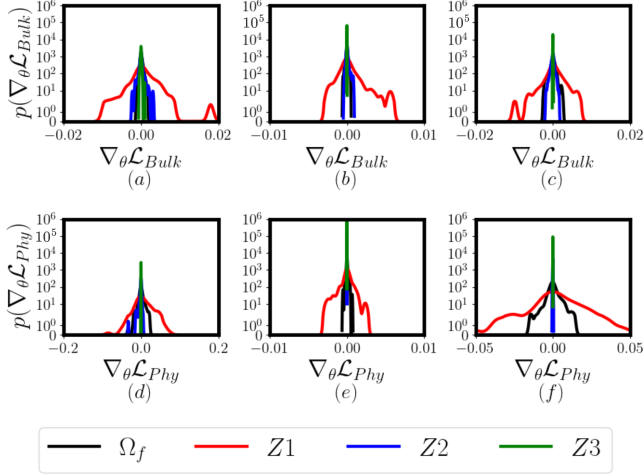


Figure 6: First layer $\nabla_{\theta}\mathcal{L}_{Bulk}$ (top row) and $\nabla_{\theta}\mathcal{L}_{Phy}$ (bottom row) distributions obtained after $1e06$ training iterations for (a,d) Case 1: $\lambda_{Phy} = 1$, (b,e) Case 2: $\lambda_{Phy} = 0.001$ and $S_{\omega_z} = 100\%$ and (c,f) Case 3: $\lambda_{Phy} = 0.01$ and $S_{\omega_z} = 5\%$.

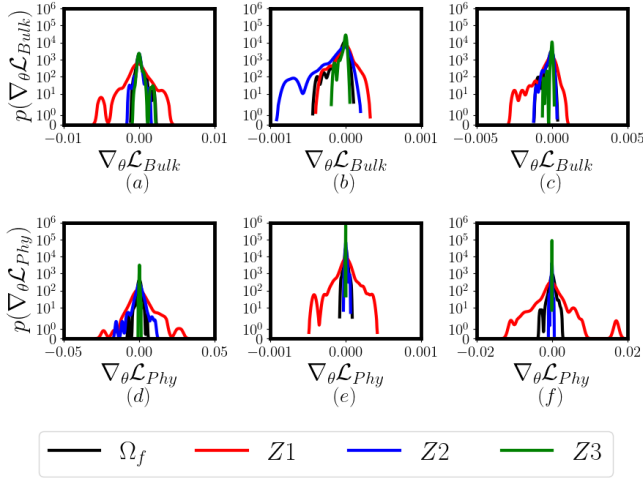


Figure 7: Comparison of last layer $\nabla_{\theta}\mathcal{L}_{Bulk}$ (top row) and $\nabla_{\theta}\mathcal{L}_{Phy}$ (bottom row) distributions obtained after $1e06$ training iterations for (a,d) Case 1: $\lambda_{Phy} = 1$, (b,e) Case 2: $\lambda_{Phy} = 0.001$ and $S_{\omega_z} = 100\%$ and (c,f) Case 3: $\lambda_{Phy} = 0.01$ and $S_{\omega_z} = 5\%$.

or even purely data-driven methods to determine which input subdomain plays a relatively active role in training.

ACKNOWLEDGEMENTS

The authors would like to sincerely thank the P.G. Senapathy High Performance Computing Centre of IIT Madras for providing the computing resources used in this study.

NOMENCLATURE

Re	Reynolds number
\mathcal{L}	Loss function
$\lambda_{\#}$	Individual loss component weights
Z	Moving body zone
$Z2$	Wake zone
$Z3$	Outer zone
$\mu_*^{Z^i}$	Zonal (Z^i) proportion weighted relative mean gradient magnitudes of loss components (*)
$\sigma_*^{Z^i}$	Zonal (Z^i) proportion weighted relative standard deviation of gradient magnitudes of loss components (*)

REFERENCES

- [1] 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*, vol. 378, pp. 686–707, 2019.
- [2] C. S. Peskin, "The immersed boundary method," *Acta numerica*, vol. 11, pp. 479–517, 2002.
- [3] M. Balajewicz and C. Farhat, "Reduction of nonlinear embedded boundary models for problems with evolving interfaces," *Journal of Computational Physics*, vol. 274, pp. 489–504, 2014.
- [4] R. Sundar, D. Majumdar, D. Lucor, and S. Sarkar, "Physics-informed neural networks modeling for systems with moving immersed boundaries: application to an unsteady flow past a plunging foil," 2023.
- [5] D. Majumdar, C. Bose, and S. Sarkar, "Capturing the dynamical transitions in the flow-field of a flapping foil using immersed boundary method," *Journal of Fluids and Structures*, vol. 95, p. 102999, 2020.
- [6] S. Wang, Y. Teng, and P. Perdikaris, "Understanding and mitigating gradient flow pathologies in physics-informed neural networks," *SIAM Journal on Scientific Computing*, vol. 43, no. 5, pp. A3055–A3081, 2021.
- [7] Z. Xiang, W. Peng, X. Liu, and W. Yao, "Self-adaptive loss balanced physics-informed neural networks," *Neurocomputing*, vol. 496, pp. 11–34, 2022.
- [8] S. J. Anagnostopoulos, J. D. Toscano, N. Stergiopoulos, and G. E. Karniadakis, "Residual-based attention and connection to information bottleneck theory in pinns," *arXiv preprint arXiv:2307.00379*, 2023.
- [9] G. C. Lewin and H. Haj-Hariri, "Modelling thrust generation of a two-dimensional heaving airfoil in a viscous flow," *Journal of Fluid Mechanics*, vol. 492, p. 339, 2003.
- [10] M. S. U. Khalid, I. Akhtar, H. Dong, N. Ahsan, X. Jiang, and B. Wu, "Bifurcations and route to chaos for flow over an oscillating airfoil," *Journal of Fluids and Structures*, vol. 80, pp. 262–274, 2018.
- [11] D. P. Kingma and J. Ba, "ADAM: A method for stochastic optimization," 2014.
- [12] A. B. Buhendwa, S. Adami, and N. A. Adams, "Inferring incompressible two-phase flow fields from the interface motion using physics-informed neural networks," *Machine Learning with Applications*, vol. 4, p. 100029, 2021.
- [13] D. Lucor, A. Agrawal, and A. Sergent, "Simple computational strategies for more effective physics-informed neural networks modeling of turbulent natural convection," *Journal of Computational Physics*, vol. 456, p. 111022, 2022.