POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

Executive Summary of the Thesis

Aortic wall shear stress quantification from 4D flow data with physics informed neural networks

Laurea Magistrale in Mathematical Engineering - Ingegneria Matematica

Author: Francesco Songia

Advisor: Dr. Stefano Pagani

Co-advisors: Dr. Francesco Regazzoni, Dr. Simone Saitta

Academic year: 2022-2023

1. Introduction

Over the past three decades, near-wall hemodynamics has been the subject of several cardiovascular fluid mechanics research activities aimed at accurately estimating biomarkers, such as wall shear stress (WSS), useful for characterizing and assessing risk in cardiovascular diseases. 4D flow magnetic resonance imaging (MRI) could be exploited to quantify WSS, providing time-resolved measurements of blood flow velocity fields in a volume of interest. However, these measurements suffer from low spatio-temporal resolution and are affected by noise. The objective of this thesis is to develop a deep learning framework based on physics informed neural networks (PINNs) to regularize and denoise 4D flow MRI recordings of blood velocity in the ascending aorta. Physics-based differential models can be considered as another source of information, as they describe the flow evolution from a modelling point of view. This balance, between available data and theoretical knowledge, can be controlled by PINNs that encode differential models into suitable terms of the loss function ensuring that the predicted velocity and pressure fields follow the underlying physical laws. Training a PINN is very challenging since the loss has several terms and the fitting data do not provide clear information as they could be corrupted and noisy. The main goal of this thesis is thus to develop an optimal configuration to present regular and denoised velocity and pressure fields. In addition, starting from the model prediction, it is accurately recovered WSS, a quantity particularly sensible to approximation errors. WSS is estimated through different methods trying to recover realistic gradients of the velocity profile near the wall. This stands as the most challenging aspect encountered, as PINNs suffer in representing high frequencies and they propose flat profiles at wall boundaries. In Figure 1, a graphical summary of this thesis is presented.

2. Problem setting

The approach developed in this thesis is based on PINNs, firstly introduced by Raissi et al in 2018 [2]. They are constructed as multi-layer perceptrons (MLPs), which receive spatio-temporal coordinates as input and reconstruct physical quantities of interest, such as velocity and pressure fields in the computational fluid dynamics case. Following [4], a sinusoidal activation function is used and the weights

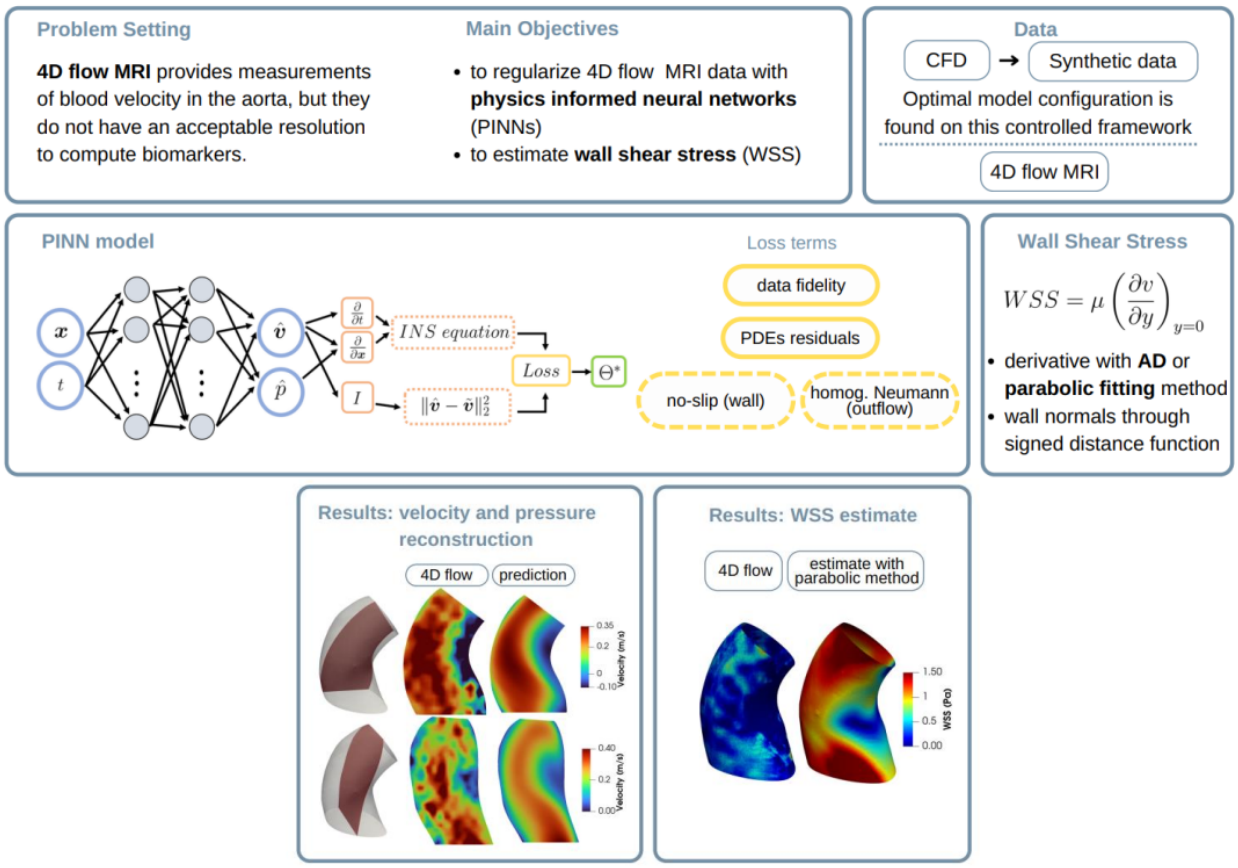

Figure 1: Overview of the thesis. After defining the motivation and the main objectives, there are represented the model structure and few results for velocity reconstruction and WSS estimate.

are uniformly initialized. The loss function is designed to balance available 4D flow MRI velocity measurements with physical knownledge encoded through the residuals of the incompressible time dependent Navier-Stokes (NS) equations. To built those residuals, velocity and pressure gradients are exactly computed through automatic differentation (AD). About boundary conditions, with additional terms in the loss, we enforce the homogeneous Neumann condition in the outflow boundary and the no-slip condition in the wall boundaries.
The global loss is a weighted sum of the different terms. The data-fidelity term is the most relevant one, while physical terms are simply kept low with a slow decreasing behaviour. Since the main goal is to denoise and regularize corrupted measurements, we avoid the overfitting of the noise thanks to the physical terms, that penalize unphysical behaviour. Nevertheless, weighting more the residual term, might end up with a flatter output. Indeed, details are challenging to recover from the PDEs and a constant zero output is always favoured by the physical terms. Provided an optimal balance, physical laws are correctly imposed and they are effective: pressure is reconstructed thanks solely to the NS momentum equation and it is given effortlessly together with the regularization process of velocity data.

No-slip boundary conditions are enforced on the wall boundary of the domain. They could be imposed with a further loss term enforcing null velocity on a subset of points or by applying the lifting procedure of CFD simulations. The latter proposes to multiply only the velocity output by a function $\Phi(\mathbf{x}) : \Omega \to \mathbb{R}$ that is null at the wall boundary $\partial\Omega$ and it increases when moving away from it. This method could be considered as it lightens the loss from one further term. The lifting function Φ can be computed analytically [1] only for simple geometries, such as a 2D toy problem. When the domain is more complex, as in 3D cases, Φ is learned by a further neural network and it represents the distance function

from the wall. In this thesis, we employ the lifting procedure only for the 2D test case, while in synthetic and real 3D test cases, the no-slip condition is imposed with a further loss term. Despite the potential advantages, using the lifting function in 3D tests does not succeed in predicting a zero wall velocity. This is because the neural net representing the wall distance in 3D is not exactly null at the boundary and there is nothing that penalizes a non-zero wall velocity.

3. Wall Shear Stress analysis

Wall shear stress is defined as

$$WSS = \mu \left(\frac{\partial v}{\partial y} \right)_{y=0},$$

where μ is the dynamic viscosity, v is the component of the velocity vector that is locally parallel to the wall, and y is the Euclidean distance from the wall. This derivative can be computed following three approaches:

1. exploiting automatic differentiation on the predicted velocity;
2. parabolic fitting method (PAR): for each wall points p_i with normal $\boldsymbol{n_i}$, this method considers two further inner points on the normal direction spaced out by d_n. The net is evaluated in those three points, then tangential velocity is computed and, with these values, a parabolic profile is fitted. The derivative is finally computed analytically;
3. variation of the parabolic fitting method with null velocity imposed on wall points (PAR ZERO).

In the parabolic fitting methods (2, 3) discrete normals coming from the mesh structure are used, while in the first case, a continuous distance function is evaluated to obtain the normals. About the latter, a neural network learns the signed distance function ϕ from the wall, then, normals are computed as $|\nabla\phi|/\|\nabla\phi\|$, where gradients are calculated through AD.
Finally, the choice of d_n influences the estimate: with a small value a very steep and variable profile is fitted and this noisy behaviour is something that is not expected. Instead, taking a larger value turns out to be the most robust choice, as it permits to rely less on the prediction near the boundary and, in any case, high velocity gradients are recovered.

Reference WSS fields. For velocity and pressure a CFD simulation represents the ground-truth when dealing with synthetic test cases. However, calculating WSS, which is a derived quantity, can result in non-negligible approximation errors unless using very fine grids. To minimize this error, the WSS is post-processed with AD on a fine neural network interpolation of the velocity. We specifically design a test case to study the accuracy in the WSS reconstruction: a small cube embedded in the aorta is taken into account and an interpolation-only neural network is trained with clean CFD velocity data. Since the domain is very small, the net well represents the velocity in the cube. Consequently, AD computes an accurate WSS field in the wall boundary and this is a first WSS reference. Moreover, those values are used to tune the optimal d_n to use within the parabolic fitting method. A larger d_n, equal to $7 \cdot 10^{-4}$ m, is found, and it used to compute a second reference WSS field in the whole geometry starting from CFD velocity data.

4. Synthetic test cases

4.1. Velocity and pressure reconstruction

Starting from the CFD reference ground-truth, three different synthetic test cases are generated with a different level of noise added (mild, medium, extreme). The procedure to create them emulates a realistic 4D flow acquisition. Data are temporally downsampled with a moving average to obtain measurements each 40 ms; then the acquisition degradation consists in a coarse frequency sampling in the corresponding k-space domain and in the addition of Gaussian noise [3]. The normalized root mean squared errors (NRMSE) between the created data and the CFD solution, are reported in Table 1, and the synthetic cases are visually compared in Figure 2.

	Vel. X	**Vel. Y**	**Vel. Z**
mild	0.013	0.013	0.014
medium	0.081	0.059	0.101
extreme	0.119	0.114	0.122

Table 1: NRMSE for different synthetic test cases.

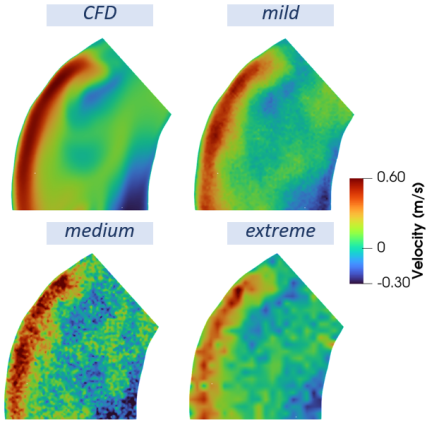

Figure 2: Velocity Z component colormaps in a 2D sagitally oriented slice.

A specific model is trained on each synthetic case. Slightly better performances are reported for the cleaner mild data, but in all the test cases, PINNs succeed in proposing super-resolved and denoised velocity and pressure fields. In the extreme case, fitting data has a velocity NRMSE approximately of 12% with respect to the CFD solution, and the model improves the representation reaching an error of 8%. Moreover, the errors with respect to the input synthetic data increase with the level of noise considered. This underlines the fact that the model prioritizes the physical regularization to the fitting data when noise is large, even if the weight of the data-fidelity term is the same. In Figure 3, 4 there are the results for the extreme case visualized in 2D slices of the aortic geometry considered.

4.2. Wall Shear Stress estimation

WSS is now quantified starting from the predicted velocity near the wall boundary. In this region, the model suffers in representing realistic velocity gradients yielding to a challenging WSS estimate. Hence, before applying the described methods to compute WSS, we try to improve the quality of the velocity representation at the wall by proposing two changes to the baseline model:

1. more collocations points are sampled in the near wall region to better enforce the governing PDEs;
2. the relevance of the loss term, that imposes a null wall velocity, is increased in order to obtain a steeper profile.

However, both these strategies does not improve

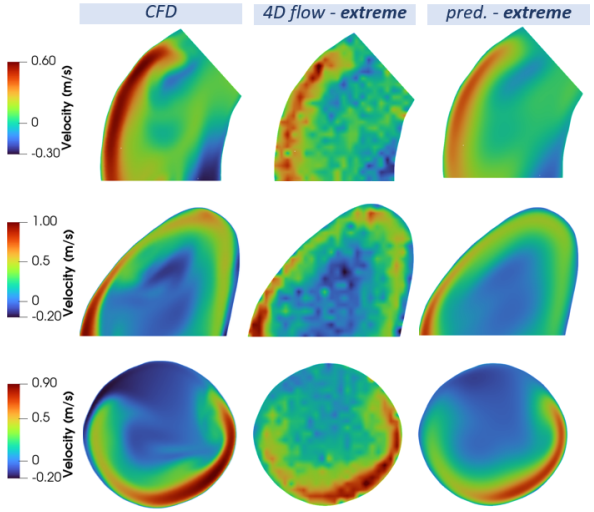

Figure 3: Velocity Z component colormaps in 2D slices. There are the CFD ground-truth, the fitting synthetic data and the model prediction for the extreme synthetic case.

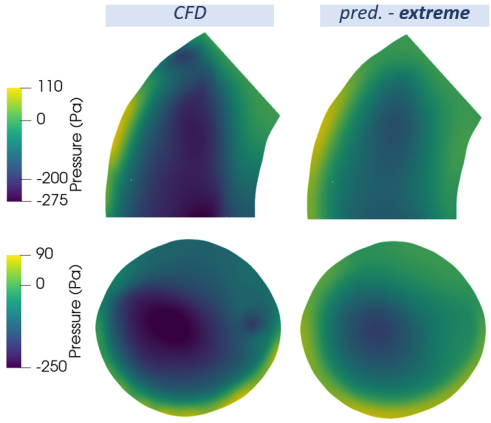

Figure 4: CFD pressure and model prediction for the extreme synthetic case.

the performances. Implicitly relying more on physics, with more collocation points, ends up in flatter profiles, as a constant velocity is favored by the loss. About the no-slip condition, increasing the loss weight related to null wall velocity succeeds in predicting a lower value at the boundary, but it does not improve the WSS estimate. This is because the velocity profile arrive at the wall with a smoother profile and, in some cases, the concavity is inverted. The initial idea was to force more the no-slip condition to end up with a steeper profile, but the results show the opposite behaviour, yielding to severe errors. Hence, the baseline model remains the best one.

To quantify WSS, we exploit the three methods to compute the normal derivative described in Section 3. AD does not provide a good estimate as it is extremely related to the poor net performances near the boundary. The parabolic fitting methods, instead, are more reliable even if they super-impose a profile that cannot always represent the reality. In particular, using a parabola that starts with a null velocity (PAR ZERO) permits to recreate the typical high velocity gradients at the wall. Moreover, it is used a large d_n to avoid relying more on the prediction near the wall. As notation, PAR2 and PAR7 correspond to the parabolic method when d_n is equal to $2 \cdot 10^{-4}$m and $7 \cdot 10^{-4}$m, respectively. The errors for the mild case are reported in Table 2.

AD	PAR2	PAR7	PAR7 ZERO
0.28	0.27	0.20	0.14

Table 2: NRMSE for WSS estimated with different methods. The reference WSS field is computed through the parabolic fitting method with $d_n = 7 \cdot 10^{-4}$m starting from CFD data.

Summing up, all methods underestimate WSS, since the model suffers in representing the near wall region. However, by applying the parabolic profile that imposes a null velocity at the wall, we obtain higher and realistic values. PAR7 ZERO is the most robust choice, and it is represented in Figure 5 where it can be visualized the improvement with respect to the WSS estimate with unprocessed 4D flow data.

5. Real 4D flow MRI test cases

We consider four measurements on different patients with bicuspid (BAV) and tricuspid (TAV) aortic valves. In particular, BAV patients are characterized by evident alterations of WSS distribution and peak values in the ascending aorta, thus, an accurate WSS estimate could provide an added value in a risk stratification study about aortopathies.
We apply the optimal PINN configuration, that is found within the study on synthetic test cases, on anonymized data acquired by Weill Cornell Medicine (NY, USA) to predict velocity and pressure, and to estimate WSS.

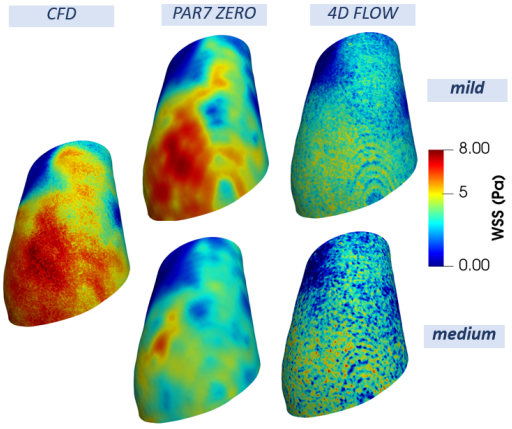

Figure 5: On the left there is the reference WSS computed with the CFD solution. In the middle column there is the estimate obtained through PAR7 ZERO, and, in the last one, there is the result obtained from raw synthetic 4D flow data.

PINNs remarkably improve the starting blood flow velocity measurements: the output fields are super-resolute and regular, as can be visualized in Figure 6, referring to a BAV patient. Moreover, physical regularization provides a realistic pressure reconstruction that can be compared with the one obtained in synthetic test cases. To support the effect of the physical regularization, we verify the mass conservation principle on a small cube inside the domain by computing the outward surface velocity flux. PINNs succeed in enforcing the mass conservation principle, indeed the computed value is approximately two orders of magnitude lower w.r.t unprocessed 4D flow measurements.
The same issues encountered with the previous test cases are detected. The neural network is not able to recover high frequencies and details are lost; in addition, the near wall region is not completely understood and the no-slip condition is not always correctly enforced. This is expected, since real data are more damaged near the boundaries due to inaccuracies in the acquisition and segmentation procedures.
In Figure 7, the results for WSS are reported for the same BAV patient. Here, all the estimates are probably underestimated, but there is a significant improvement with respect to unprocessed 4D flow data. Only the parabolic fitting method considering a parabola that starts from zero succeeds in proposing higher values.

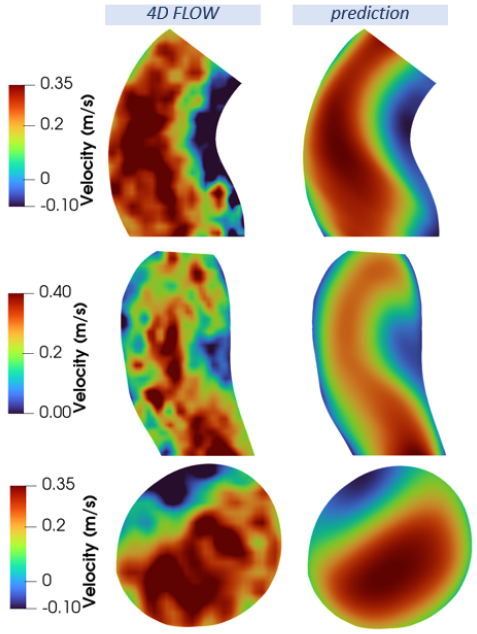

Figure 6: Velocity Z component colormaps in 2D slices. The model is trained with data on the left column, and, on the right, there are its super-resolute predictions.

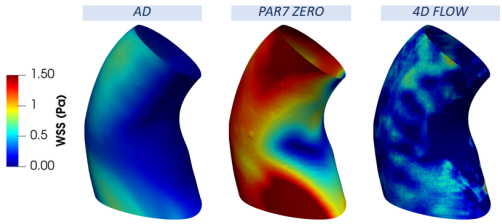

Figure 7: WSS computed with AD and with the parabolic fitting method with null velocity for the wall points, and $d_n = 7 \cdot 10^{-4}$ m.

6. Conclusion

Blood flow measurements in the aorta provided by MRI techniques are affected by significant noise sources, making them unsuitable for investigating biomarkers such as WSS. Nevertheless, a well known modelling knowledge could describe the blood flow evolution and this another source of information is used to regularize and denoise the initial measurements. Our results show how PINNs improve the quality of those velocity and pressure reconstructions and how they permit to provide an estimate of WSS. Moreover, we propose an alternative method to compute wall normals starting from a neural network that represents the signed distance function from the boundary.

The reconstruction of realistic velocity profile near the boundary, that mainly characterize the stress against the vessels, remains very challenging. The proposed neural networks suffer from representing high frequencies and details. Indeed, there is not visible the flow complexity typical of realistic 3D CFD simulations. All the predictions show a smooth behaviour and this leads to an underestimation of WSS. To face this issue, we pay particular attention to the near wall region by adding more collocation points where evaluating the PDEs residuals, and by trying to recover the usual high velocity gradients.

Finally, the computational resources required make this methodology not yet feasible in real applications, as they require results in a very concise time. The net training is heavy, and, even if the model structure and the weights used in the loss are designed to remain unchanged, the training must be patient-specific. Further developments could go in the direction of exploiting transfer learning techniques to reuse the information discovered when analyzing a specific patient.

References

[1] S Berrone et al. Enforcing dirichlet boundary conditions in physics-informed neural networks and variational physics-informed neural networks. *arXiv preprint arXiv:2210.14795*, 2022.

[2] M Raissi et al. 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.

[3] S Saitta et al. Implicit neural representations for unsupervised super-resolution and denoising of 4d flow mri. *arXiv preprint arXiv:2302.12835*, 2023.

[4] V Sitzmann et al. Implicit neural representations with periodic activation functions. *Advances in neural information processing systems*, 33:7462–7473, 2020.