Hemodynamic Neural Field: Continuous Blood Flow Estimation in Vessel Geometries via Global Shape Conditioning

Simon Perrin[1][0009−0005−2008−2802], Patryk Rygiel[2][0009−0003−8539−5581], Amin Ranem[2][0000−0003−0783−6903], Bram de Wilde[2,3][0000−0003−1890−8714], Sébastien Levilly[1][0000−0002−2704−0315], Jelmer M. Wolterink[2][0000−0001−5505−475X], and Harold Mouchère[1][0000−0001−6220−7216]

[1] Nantes Université, École Centrale Nantes, CNRS, LS2N, UMR 6004, F-44000 Nantes, France
simon.perrin@univ-nantes.fr
[2] Department of Applied Mathematics, Technical Medical (TechMed) Centre, University of Twente, Enschede, The Netherlands
[3] Physics of Fluids Group, Technical Medical (TechMed) Centre, University of Twente, Enschede, The Netherlands

Abstract. Hemodynamic biomarker estimation is crucial for the diagnosis and prognosis of cardiovascular diseases. Patient-specific computational fluid dynamics (CFD) simulations offer high-fidelity data but are computationally expensive and time-consuming, limiting their immediate implementation in resource-constrained clinical settings. To address this challenge, we propose Hemo-NF, a two-stage conditional neural field framework for continuous blood flow prediction in vessel geometries. Unlike state-of-the-art architectures whose complexity scales with explicit mesh size, Hemo-NF uses global shape conditioning to achieve constant processing complexity. It encodes vessel geometries using DeepSDF via an auto-decoding strategy and conditions a secondary neural field to predict three-dimensional blood flow velocity components. To bridge the gap between global shape representation and local geometric one, we propose to use locally extracted features from DeepSDF to condition the flow prediction network. Validated on a synthetic dataset of bifurcating arteries, Hemo-NF achieves competitive results compared to SEGNN when considering successful geometry encodings. We also provide a comprehensive evaluation of different conditioning mechanisms (concatenation vs. FiLM) and latent representations. This framework offers a promising approach to continuous, geometry-driven hemodynamic assessment.

Keywords: Blood Flow Prediction · Conditional Neural Fields · Global Shape Conditioning

1 Introduction

In the context of cardiovascular diseases, hemodynamic biomarker estimation allows diagnosis and pronosis [16, 5]. Patient-specific Computational Fluid Dy-

namics (CFD) based on the vessel's geometry and hemodynamic boundary conditions provides a high-fidelity and high-resolution simulation, enabling highly accurate estimation of biomarkers. Nonetheless, the high computational cost and long processing time can make this solution challenging to implement in clinical settings with limited resources. This highlights the need for surrogate methods that can estimate CFD results rapidly, reliably, and with limited resources.

Deep learning approaches allow blood flow prediction in a vessel geometry. This type of solution allows to predict velocity components in few seconds with a relative low error rate. Convolutional Neural Network (CNN) solutions have been used to predict blood flow in stenotic arteries [15] or estimate Wall Shear Stress [3]. Nevertheless, CNN-based approaches are constrained by the resolution of the underlying cartesian grid. This limitation restricts their capacity to capture complex vascular geometries and fine structural details, which critically govern local hemodynamics and highly sensitive blood flow patterns. To overcome this, point-cloud and mesh-based solutions have been proposed to directly handle complex shape (e.g. LaB-GATr [13]). Furthermore, Graph Neural Networks (GNN)-based architectures are well-suited for modeling topological features and capturing local dynamics through message-passing protocols (e.g. SEGNN [11]). These geometric deep learning approaches can be broadly categorised into unreduced modeling, which operates directly on explicit full-resolution mesh vertices (e.g. SEGNN [11]), and reduced modeling, which utilises compressed or tokenized latent spaces (e.g. LaB-GATr [13]). Furthermore, these methods require the entire mesh to be processed, which means that their complexity scales with the mesh size. Neural Fields have emerged as a continuous framework to model physical phenomena. Physics-Informed Neural Networks, which are a special case of neural fields using physics-based losses, have been used to perform super-resolution of 4D Flow MRI [2, 10] or estimate arterial blood pressure [4]. Although these solutions can provide high-fidelity results, they cannot generalise to new geometries and must be retrained for each new patient, limiting clinical utility. To avoid retraining for each new data, neural fields can be conditioned with a latent vector to generalise on a task, using methods such as concatenation with spatial coordinates [6] or FiLM modulations [7]. Conditional Neural Fields (CNFs) have been used for turbulence prediction [1] or flow prediction in a specific geometry [9].

In this work, we propose Hemo-NF, a CNF for blood flow prediction in vessel geometries. While Hemo-NF positions itself as a reduced modeling framework, unlike methods like LaB-GATr [13] that rely on multiple localised sparse tokens, we address the question of **using a global latent representation of the shape to capture the vessel geometry for accurate continuous hemodynamic prediction**, achieving a constant complexity w.r.t. mesh size. The vessel geometry is globally encoded via gradient descent on DeepSDF [6], a geometry CNF, yielding a latent code that conditions a second neural field for flow prediction. Moreover, to bridge the gap between global shape understanding and local geometrical details, we propose to condition the second neural field on local information extracted by DeepSDF. Hemo-NF has been trained

and validated on a dataset of synthetic bifurcating arteries with ground truth from CFD. We find that Hemo-NF outperforms PointNet++ [8], and close in performance to SEGNN [11]. Although a performance gap remains with state-of-the-art local methods such as Lab-GATr [13], our results demonstrate that a global continuous representation of the shape geometry can effectively enable blood flow prediction while achieving constant processing complexity w.r.t. mesh size.

2 Method

2.1 Data

To evaluate and train Hemo-NF for hemodynamic prediction, we use the public synthetic dataset introduced in [12]. The dataset is composed of 2 000 bifurcating arteries synthesized after measurement statistics of left main coronary bifurcations. The dataset is split as 1600:200:200 between training, validation and evaluation. Each artery geometry is composed of $\sim 175,000$ vertices ($\sim 17,000$ on surface). Steady-flow CFD simulation for each geometry has been performed using SimVascular [14]. The same initial and boundary conditions are used for all geometries, enabling blood flow to be predicted solely based on the geometry. A low Reynolds number was used, resulting in almost laminar flow. More simulation details can be found in [12].

2.2 Hemo-NF

Hemodynamic Neural Field (Hemo-NF) comprises two distinct neural fields: the first acts as an auto-decoder to learn a latent representation of the shape geometry, which then conditions the second neural field to predict blood flow. A schematic representation of the network architecture is shown in Fig. 1.

Geometry encoding using DeepSDF To perform the vessel geometry encoding, we use the DeepSDF [6] architecture (f). DeepSDF represents geometry by learning a Signed Distance Function (SDF), which maps spatial coordinates to their signed distance from the surface boundary. The function's output represents the distance magnitude, while the sign indicates whether a point lies inside (negative) or outside (positive) the shape. Then, the underlying surface geometry can be extracted by taking the zero-level set of the SDF. The network architecture is a standard CNF that is implemented using a multi-layer perceptron (MLP) with ReLU activation functions. Tanh is used as the final activation function, and Dropout layers are used after each hidden layer. It takes the concatenation of three-dimensional spatial coordinates and the latent representation of the vessel's shape (z_g) as input and predicts the SDF value. As DeepSDF is using an autodecoding learning strategy, the latent vector z_g is obtained with gradient descent on the neural network. Specifically, for each new geometry, the global latent representation z_g is learned by performing gradient descent starting from a zero vector.

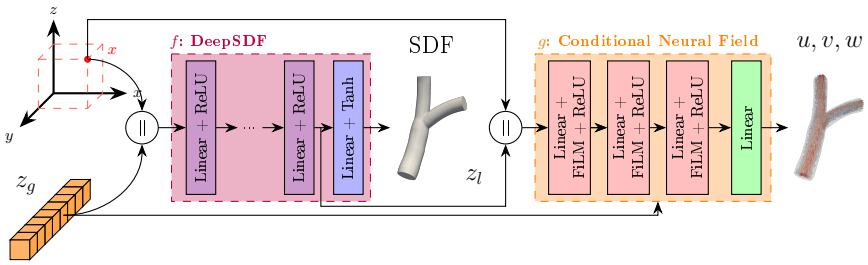

Fig. 1: Hemo-NF architecture. DeepSDF (f) is used to learn the global latent representation z_g and extract the local representation z_l. Meanwhile, the secondary conditional neural field (g) predicts the three blood flow velocity components (u, v, w), conditioned on both z_g and z_l. The symbol || means concatenation.

Blood flow prediction A second conditional neural field (g), implemented as an MLP using ReLU activation functions, predicts the 3D blood flow velocity vector (u, v, w). Like DeepSDF, it is conditioned on the global geometric latent vector z_g. To enhance conditioning with local information, we also introduce the local conditioning z_l such as $z_l = f_{-1}(z_g, x)$, where f_{-1} represents the f network truncated before its final output layer and x the 3D spatial coordinate. z_l is representing local information extracted by DeepSDF at the coordinate x. Geometric information is propagated into g via a combination of input concatenation and FiLM feature modulation [7]. First, z_l is concatenated with x at the network input. Then, z_g modulates each hidden layer via a predicted FiLM bias:

$$x_i = \sigma(W_i x_{i-1} + B_i + W_i^{\text{cond}} z_g), \tag{1}$$

where W_i, B_i are the i-th layer weights and biases, x_{i-1} the output of the previous layer, σ is the activation function, and W_i^{cond} the weights to predict the FiLM bias from z_g. This hybrid conditioning (input concatenation for z_l and FiLM for z_g) yielded the best performance in our empirical analysis (see Sec. 3.2).

2.3 Training strategy

We train Hemo-NF in two steps. Firstly, DeepSDF is trained on the geometry prediction task. The loss function is defined as:

$$\mathcal{L}_{\text{geo}} = |\text{clamp}(f(x, z_g), \delta) - \text{clamp}(s, \delta)| + \lambda ||z_g||_2^2, \tag{2}$$

where $\text{clamp}(x, \delta) = \min(\delta, \max(-\delta, x))$, δ controls the SDF truncation distance, f the DeepSDF network and $||z_g||_2^2$ is the regularisation norm of the global latent vector z_g, as in [6]. As f is trained in an autodecoding way, the conditioning vectors z_g are trained at the same time as the neural network parameters for the training data. For the test data, these vectors are optimized afterwards while keeping the network weights frozen. Secondly, the CNF for blood flow prediction is trained using a $\mathcal{L}_1$ loss between the velocity predicted and ground truth.

Table 1: Quantitative evaluation of blood flow prediction. The metrics are estimated by comparing the prediction of approaches with the CFD simulation for each geometry at every vertex. The results for PointNet++, SEGNN and LaB-GATr are taken from [11, 13].

Model	**NMAE**		**Approximation error**	
	average	median	average	median
Hemo-NF w/o z_l	1.51 %	1.28 %	13.44 %	10.91 %
Hemo-NF	**1.20 %**	**0.95 %**	**10.44 %**	**7.99 %**
PointNet++ [11]	1.1 %	n/a	11.0 %	n/a
SEGNN [11]	0.7 %	n/a	7.4 %	n/a
LaB-GATr [13]	n/a	n/a	3.3 %	n/a

3 Evaluation of Hemo-NF

3.1 Implementation

Following [6], DeepSDF comprises 8 layers of 512 neurons. The conditioning relies on a global geometric latent vector z_g of size 256 and a local feature vector z_l of size 512. DeepSDF is trained via Adam with a batch size of 16 and a learning rate of 3×10^{-5} (1×10^{-5} for z_g) during 600 epochs. Loss parameters are set to $\lambda = 0.01$ and $\delta = 0.1$ (Eq. 2). For each geometry, $32,768$ points are randomly sampled (80% near the surface, 20% globally). Latent vector z_g optimisation for new geometries uses the same learning rate over 200 epochs. Then, the CNF (g) is made of 4 layers of 512 neurons and trained via Adam during $1,000$ epochs with a batch size of 16, a learning rate of 1×10^{-4}, and $32,768$ points randomly sampled across the surface and volume for each geometry. Velocity values are normalised to $[-1, 1]$ for training and reverted for evaluation. Finally, blood flow prediction in test geometries is evaluated at each vertex, after learning z_g, using Normalised Mean Absolute Error (NMAE) and Approximation Error [11, 13].

3.2 Results

Table 1 presents quantitative results for Hemo-NF using or not z_l conditioning. Moreover, we compared our results with PointNet++ [8], SEGNN [11], and LaB-GATr [13]. Hemo-NF performed better than PointNet++ considering approximation error but not on the NMAE metric. This discrepancy suggests that PointNet++ focuses more on accurately reflecting high-velocity regions, which are heavily penalised by NMAE's normalisation, whereas Hemo-NF provides a more balanced prediction that is less biased toward these peak-flow areas. However, a performance gap remains when comparing our framework to state-of-the-art architectures like SEGNN and LaB-GATr. Finally, the gap between the average and median results of Hemo-NF reveals inaccurate predictions caused by geometry encoding failure.

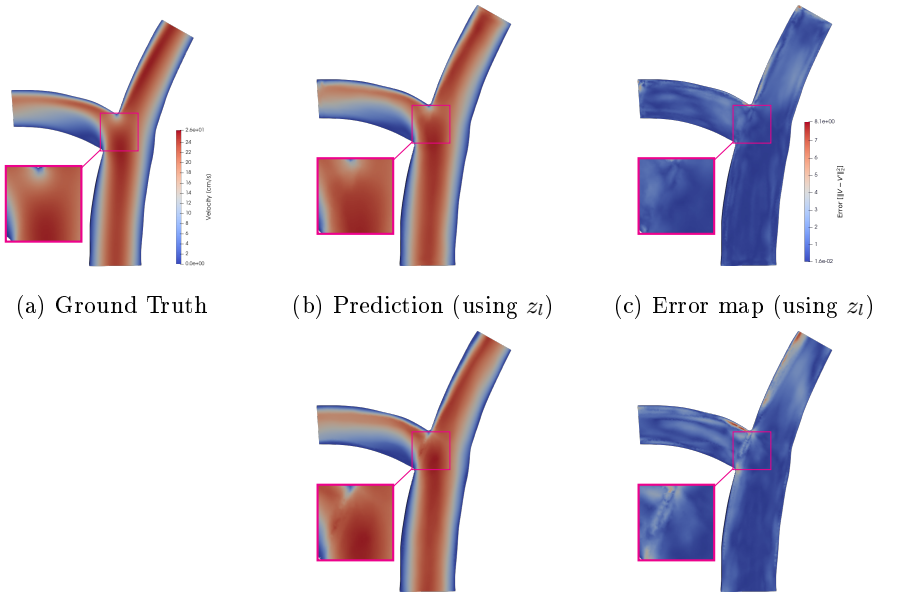

(a) Ground Truth (b) Prediction (using z_l) (c) Error map (using z_l)

(d) Prediction (w/o using z_l) (e) Error map (w/o using z_l)

Fig. 2: Visualisation of blood flow prediction on a geomery from test split. (a) CFD ground truth. (b) and (c) are prediction and error map for our approach using z_l. (d) and (e) are the same without using z_l. The zooms highlight that incorporating z_l significantly reduces prediction error at the bifurcation point.

Impact of local conditioning z_l The use of the coordinate dependent conditioning z_l improves blood flow prediction as shown by the two metrics, with a gain of 0.31 % on NMAE and 3 % for the approximation error. This conditioning gives geometrical information for each point in the space, such as the distance to the surface. Figure 2 shows qualitative results on a test geometry comparing our approach with and without using z_l. A noticeable difference can be seen at the bifurcation point, with an error line in the prediction when z_l is not used. The use of z_l also helps to significantly minimise the error close to the surface at the top of the vessel.

Conditioning strategy analysis We investigate different combinations of concatenation and FiLM modulation to inject the global z_g and local z_l geometric representations into the neural field. As shown in Table 2, solutions relying on a single conditioning strategy (#1, #2) generally underperform compared to hybrid methods. This demonstrates that combining both conditioning mechanisms improves blood flow prediction. Configurations #3 and #4 evaluate disjoint strategies, where each representation is assigned to a distinct conditioning method. Configuration #4 achieves the best overall performance (**10.44**), lever-

Table 2: Results for different ways of conditioning our solution with the coordinate-based and global latent representations.

Configuration	Concatenation		FiLM		Appr. error
	$\mathcal{Z}_g$	$\mathcal{Z}_l$	$\mathcal{Z}_g$	$\mathcal{Z}_l$	
#1	×	×			11.36 %
#2			×	×	10.97 %
#3	×			×	10.78 %
#4		×	×		**10.44** %
#5	×	×	×	×	<u>10.50</u> %

aging FiLM to modulate hidden layers with the global shape context z_g, while concatenating the coordinate-dependent features z_l at the network input to locally adjust the prediction. Interestingly, providing both representations to all mechanisms (#5) yields the second-best results (<u>10.50</u>), suggesting that the network can learn to filter redundant information.

Results depending on geometric encoding error Unlike methods that directly process the explicit mesh geometry to predict the flow [11, 13], our approach relies on learning a global shape latent code. In Fig. 3, we examine the results in relation to encoding quality. We utilise the Chamfer Distance to evaluate the quality of the geometric reconstruction derived from the encoding against the ground truth. Furthermore, we group the assessment of blood flow prediction according to quartiles of geometric reconstruction quality. That figure reveals a direct relationship between the geometry encoding quality, and particularly for the last quartile with a poor Chamfer Distance, and the velocity predictions. Conversely, no such clear correlation is evident for the first three quartiles. The Pearson correlation between the two metrics is 0.26 for the three first quartiles and 0.83 for the last one. These findings highlight that while the method can deal with different qualities of shape encoding to some degree, its performance clearly deteriorates when the encoding is too poor as reflected by Q4. Further analysis of Q4 shows these errors arise from geometries lying out-of-distribution. For these atypical shapes, the auto-decoder fails to learn a valid latent code. Thus, these failures stem from the limited capacity of the learned latent space to generalise to unseen geometries, rather than from test-time optimisation issues. Finally, when considering only the first two quartiles, the metrics for flow prediction quality are close to those of the SEGNN method. This demonstrates that utilising our approach to predict flow yields results similar to those of SEGNN, and that the real bottleneck lies in the encoding of the geometry.

4 Discussion

The proposed framework's strength lies in its ability to build a compact continuous representation of both the geometry and the velocity. In contrast to

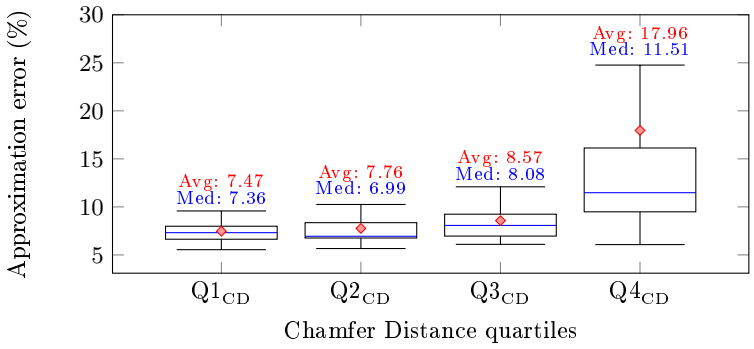

Fig. 3: Hemo-NF approximation error for every Chamfer Distance (CD) quartiles. The results highlight that the blood flow prediction quality is heavily dependent on the geometric encoding stage, whereas the poorest blood flow prediction happens on the last CD quartile. Conversely, the first two quartiles maintain low, stable errors, achieving a performance level close to SEGNN. Outliers are not displayed for clarity.

LaB-GATr [13] and SEGNN [11], our method does not scale with the size of the input mesh. Furthermore, unlike SEGNN [11], which is restricted to discrete mesh points, Hemo-NF provides a continuous representation, allowing for blood flow prediction at any arbitrary position in space. Moreover, this continuous representation using a neural field could enable a more accurate assessment of hemodynamic biomarkers dependent on spatial derivatives, such as wall shear stress (WSS). We could use auto-derivation directly on the continuous field to estimate them, avoiding the resolution limitations of the input mesh. However, while this approach avoids resolution dependencies and allows continuous prediction, a performance gap still remains with state-of-the-art methods like LaB-GATr [13], which achieves an approximation error of 3.3 %. This gap mainly comes from the chosen representation, as relying only on a single global encoding is limited and not sufficient to model highly local behaviors. Indeed, by adding the local conditioning, which depends on the global information, we successfully improved the results, reducing the approximation error from 13.44 % without local context to 10.44 % with it. This improvement highlights that purely global representation needs to be enriched with local context. Lastly, a major difference between Hemo-NF, SEGNN, and LaB-GATr lies in rotation and translation equivariance. In our method, the solution depends on the geometry's position and rotation, with this spatial data being embedded into the learned latent code. Enforcing equivariance in our approach could further enhance estimation accuracy by maximising the geometric content stored in the latent code.

5 Conclusion and Perspectives

In this work, we propose Hemo-NF, a conditional neural field approach for blood flow prediction in vessel geometry. Unlike traditional methods that depend on mesh resolution, our approach predicts the blood flow as a continuous function mapping spatial coordinates and blood flow velocities. By enriching the global geometry context with a more local representation depending on it, our approach successfully predicts blood flow with performance close to SEGNN. However, a performance gap still remains compared to LaB-GATr. Finally, we investigate the simultaneous use of two conditioning strategies to improve predictions.

The proposed strategy predicts steady flow within a blood vessel. This scenario does not fully reflect the cardiovascular system, where flow is predominantly pulsatile, particularly near the heart. As a future work, integrating the temporal dimension and pulsatile flow dynamics should be investigated. Moreover, the solution should integrate and be train on varying boundary conditions to enable patient-specific prediction aligned with clinical reality. Finally, a comprehensive clinical study on biomarker assessment derived from our continuous predictions should be conducted to validate the clinical utility of the solution.

Acknowledgments. This work was performed using HPC resources from GENCI–IDRIS (Grant 2025-AD011017146) and resources of the GLiCID Computing Facility (Ligerien Group for Intensive Distributed Computing, https://doi.org/10.60487/glicid, Pays de la Loire, France).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Du, P., Parikh, M.H., Fan, X., Liu, X.Y., Wang, J.X.: Conditional neural field latent diffusion model for generating spatiotemporal turbulence **15**(1), 10416. https://doi.org/10.1038/s41467-024-54712-1
2. Fathi, M.F., Perez-Raya, I., Baghaie, A., Berg, P., Janiga, G., Arzani, A., D'Souza, R.M.: Super-resolution and denoising of 4D-Flow MRI using physics-Informed deep neural nets **197**, 105729. https://doi.org/10.1016/j.cmpb.2020.105729
3. Kazemi, A., Stoddard, M., Amini, A.A.: Tke-net: Deep learning for estimation of super-resolved turbulent kinetic energy maps from 4d-flow mri data. In: 2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI). pp. 1–4 (2023). https://doi.org/10.1109/ISBI53787.2023.10230629
4. Kissas, G., Yang, Y., Hwuang, E., Witschey, W.R., Detre, J.A., Perdikaris, P.: Machine learning in cardiovascular flows modeling: Predicting arterial blood pressure from non-invasive 4D flow MRI data using physics-informed neural networks **358**, 112623. https://doi.org/10.1016/j.cma.2019.112623
5. Netala, V.R., Hou, T., Wang, Y., Zhang, Z., Teertam, S.K.: Cardiovascular Biomarkers: Tools for Precision Diagnosis and Prognosis. International journal of molecular sciences **26**(7), 3218. https://doi.org/10.3390/ijms26073218

6. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: Deepsdf: Learning continuous signed distance functions for shape representation. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 165–174 (2019). https://doi.org/10.1109/CVPR.2019.00025
7. Perez, E., Strub, F., de Vries, H., Dumoulin, V., Courville, A.: Film: Visual reasoning with a general conditioning layer. Proceedings of the AAAI Conference on Artificial Intelligence **32**(1) (Apr 2018). https://doi.org/10.1609/aaai.v32i1.11671
8. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: deep hierarchical feature learning on point sets in a metric space. In: Proceedings of the 31st International Conference on Neural Information Processing Systems. pp. 5105–5114. NIPS'17, Curran Associates Inc., Red Hook, NY, USA (2017)
9. Serrano, L., Boudec, L.L., Koupaï, A.K., Wang, T.X., Yin, Y., Vittaut, J.N., Gallinari, P.: Operator learning with neural fields: Tackling pdes on general geometries. 37th Conference on Neural Information Processing Systems (NeurIPS 2023) (2023). https://doi.org/10.52202/075280-3093
10. Shone, F., Ravikumar, N., Lassila, T., MacRaild, M., Wang, Y., Taylor, Z.A., Jimack, P., Dall'Armellina, E., Frangi, A.F.: Deep Physics-Informed Super-Resolution of Cardiac 4D-Flow MRI. In: Information Processing in Medical Imaging, vol. 13939, pp. 511–522. Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-34048-2_39
11. Suk, J., Brune, C., Wolterink, J.M.: SE(3) symmetry lets graph neural networks learn arterial velocity estimation from small datasets. In: Functional Imaging and Modeling of the Heart. pp. 445–454. Springer Nature Switzerland, Cham (2023). https://doi.org/10.1007/978-3-031-35302-4_46
12. Suk, J., Haan, P.d., Lippe, P., Brune, C., Wolterink, J.M.: Mesh convolutional neural networks for wall shear stress estimation in 3d artery models. In: Statistical Atlases and Computational Models of the Heart. Multi-Disease, Multi-View, and Multi-Center Right Ventricular Segmentation in Cardiac MRI Challenge. pp. 93–102. Springer International Publishing, Cham (2022). https://doi.org/10.1007/978-3-030-93722-5_11
13. Suk, J., Imre, B., Wolterink, J.M.: Lab-gatr: Geometric algebra transformers for large biomedical surface and volume meshes. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2024. pp. 185–195. Springer Nature Switzerland, Cham (2024). https://doi.org/10.1007/978-3-031-72390-2_18
14. Updegrove, A., Wilson, N.M., Merkow, J., Lan, H., Marsden, A.L., Shadden, S.C.: SimVascular: An Open Source Pipeline for Cardiovascular Simulation **45**(3), 525–541. https://doi.org/10.1007/s10439-016-1762-8
15. Wang, L., Dong, D., Tian, F.B.: Fast prediction of blood flow in stenosed arteries using machine learning and immersed boundary-lattice boltzmann method. Frontiers in Physiology **Volume 13 - 2022** (2022). https://doi.org/10.3389/fphys.2022.953702
16. Wåhlin, A., Eklund, A., Malm, J.: 4D flow MRI hemodynamic biomarkers for cerebrovascular diseases. Journal of Internal Medicine **291**(2), 115–127 (2022). https://doi.org/10.1111/joim.13392