



Contents lists available at ScienceDirect

## Computer Methods and Programs in Biomedicine

journal homepage: [www.elsevier.com/locate/cmpb](http://www.elsevier.com/locate/cmpb)

# Implicit neural representations for unsupervised super-resolution and denoising of 4D flow MRI

Simone Saitta<sup>a,\*</sup>, Marcello Carioni<sup>b</sup>, Subhadip Mukherjee<sup>c</sup>, Carola-Bibiane Schönlieb<sup>d</sup>,  
Alberto Redaelli<sup>a</sup>

<sup>a</sup> Department of Electronics, Information and Bioengineering, Politecnico di Milano, Milan, Italy

<sup>b</sup> Department of Applied Mathematics, University of Twente, 7500AE Enschede, the Netherlands

<sup>c</sup> Department of Electronics & Electrical Communication Engineering, Indian Institute of Technology (IIT) Kharagpur, India

<sup>d</sup> Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge, UK

## ABSTRACT

**Background and Objective:** 4D flow magnetic resonance imaging provides time-resolved blood flow velocity measurements, but suffers from limitations in spatio-temporal resolution and noise. In this study, we investigated the use of sinusoidal representation networks (SIRENs) to improve denoising and super-resolution of velocity fields measured by 4D flow MRI in the thoracic aorta.

**Methods:** Efficient training of SIRENs in 4D was achieved by sampling voxel coordinates and enforcing the no-slip condition at the vessel wall. A set of synthetic measurements were generated from computational fluid dynamics simulations, reproducing different noise levels. The influence of SIREN architecture was systematically investigated, and the performance of our method was compared to existing approaches for 4D flow denoising and super-resolution.

**Results:** Compared to existing techniques, a SIREN with 300 neurons per layer and 20 layers achieved lower errors (up to 50% lower vector normalized root mean square error, 42% lower magnitude normalized root mean square error, and 15% lower direction error) in velocity and wall shear stress fields. Applied to real 4D flow velocity measurements in a patient-specific aortic aneurysm, our method produced denoised and super-resolved velocity fields while maintaining accurate macroscopic flow measurements.

**Conclusions:** This study demonstrates the feasibility of using SIRENs for complex blood flow velocity representation from clinical 4D flow, with quick execution and straightforward implementation.

## 1. Introduction

4-dimensional flow encoded magnetic resonance imaging (4D flow magnetic resonance imaging (MRI) or 4D flow) is the only existing non-invasive imaging technique that provides true time-resolved 3-dimensional (3D) and 3-directional blood flow velocity measurements [18]. 4D flow is based on the phase contrast MRI (PC-MRI) principle, which makes use of bipolar magnetic gradients to calculate the phase shift of moving protons. PC-MRI encodes tissue velocity  $\mathbf{v}(\mathbf{x}, t) \in \mathbb{R}^3$  at spatial location  $\mathbf{x}$  during cardiac phase  $t$  ( $1 \leq t \leq N_t$ ) according to:

$$\rho_i(\mathbf{x}, t) = \rho_0(\mathbf{x}, t) \exp\left(j\pi \frac{(\Phi \mathbf{v}(\mathbf{x}, t))_i}{VENC}\right), \quad (1)$$

where  $VENC$  is a manually set parameter determining the maximum velocity that can be recorded,  $i = 0, \dots, 3$  are the encoded velocity components, and adopting a four-point velocity encoding,  $\Phi$  is defined as:

$$\Phi = \begin{bmatrix} 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (2)$$

Hence, the measured tissue velocity component  $i$  is proportional to the phase shift of the reconstructed images  $\rho_i$  [42].

Considering  $\rho_{it} \in \mathbb{R}^{N_r \times N_c \times N_s}$  a discretized complex PC image on a Cartesian  $N_r \times N_c \times N_s$  grid corresponding to cardiac phase  $t$  and velocity component  $i$ , the reconstructed image  $\mathbf{H}(\rho_{it}) \in \mathbb{R}^{N_r \times N_c \times N_s}$  can be modeled as:

$$\mathbf{H}(\rho_{it}) = \mathbf{F}^{-1}(\mathbf{M}(\mathbf{F}(\rho_{it}) + \epsilon)), \quad (3)$$

where  $\mathbf{F}$  is the Fourier transform,  $\mathbf{M} \in \{0, 1\}^{N_r \times N_c \times N_s}$  defines the undersampling mask in  $k$ -space, and  $\epsilon \in \mathbb{C}^{N_r \times N_c \times N_s}$  is the additive complex noise. Herein, we have neglected coil sensitivity maps for simplic-

\* Corresponding author.

E-mail address: [simone.saitta@polimi.it](mailto:simone.saitta@polimi.it) (S. Saitta).

ity, a more rigorous description of the PC-MRI measurement operator can be found in [42].

Velocity fields measured by 4D flow MRI can be processed after reconstruction to quantify more complex and clinically relevant hemodynamic biomarkers such as wall shear stress (WSS) [4,40,6], relative pressure [7,31] and vortex structure [26]. However, reconstructed flow images suffer from important limitations that make them inadequate for accurate assessment of blood flow markers. These limitations mainly concern spatio-temporal resolution, velocity encoding (VENC) related signal to noise ratio (SNR) and  $k$ -space noise [16]. Overcoming 4D flow limitations by reducing noise levels and enhancing its spatio-temporal resolution, would lead to more accurate hemodynamic assessment, boosting its widespread adoption and increasing its clinical usefulness.

In this work, we present a novel application of implicit neural representations (INRs) [34] to achieve super-resolution (SR) and denoising of 4D flow MRI velocity fields. Our approach provides a continuous reconstruction of blood flow in both space and time, showing superior performance with respect to existing baseline methods for 4D flow denoising.

### 1.1. Related work

Several lines of research have been pursued to overcome the main limitations of 4D flow MRI velocity measurements and provide more accurate non-invasive hemodynamic assessment. For the sake of brevity, we will only focus on image processing approaches, but it is important to mention other strategies based on Kalman filtering [12,14], adjoint-based optimization [39], reduced order modeling-based regression [3], and physics-informed compressed sensing [17]. Most image processing approaches employ interpolation techniques to approximate flow data, incorporating some form of regularization to impart desirable characteristics to the resulting velocity fields. Within this category, we discern between conventional denoising approaches and recent neural network-based methods.

*Conventional denoising approaches* adopt signal processing tools to enhance the acquired flow MRI data after reconstruction. Busch et al. [8], achieved denoising by projecting noisy flow measurements onto a 3D space of divergence-free radial basis functions (RBFs). Nonetheless, strict enforcement of the divergence-free condition may result in over-regularized velocity fields, especially in flow regions near edges of flow (boundaries). An approach to *softly* enforce the divergence-free condition was proposed by a Ong et al. [21]. By constructing divergence-free wavelets (DFWs), the authors essentially achieve wavelet denoising, but the different choice of wavelets enables correction of divergent flow components, providing flow fields with more coherent streamlines when compared to unfiltered medical data. By directly processing images, these data-based methods are typically characterized by high computational efficiency. However, none of the proposed approaches so far has taken into account the time-dependent nature of 4D flow data. Additionally, working on structured grids (images) prevents a precise quantification of near-wall blood flow markers that are often sought-after in medical applications, such as WSS fields. In a recent study [38], 4D RBFs were used to reconstruct velocity fields measured with MRI. To make their approach computationally light, the authors combined multi-quadratic RBF interpolation with a partition of unity scheme. Since it operates on scattered data, this approach becomes efficient in high dimensions, and well-suited for enforcing Dirichlet boundary conditions on the reconstructed velocity fields and improving quantification of near-wall features.

*Neural network approaches* aim to learn approximations of complex nonlinear functions. A popular paradigm for exploiting the high expressive power of neural networks (NNs) is represented by physics informed neural networks (PINNs). In the formulation proposed in [25], PINNs are parametrized multi-layer perceptrons (MLPs) that learn a function

mapping coordinates to outputs. PINNs can be seen as continuous functions fitting sparse observations and simultaneously minimizing the residual of a partial differential equation (PDE) (in differential form) that is identified as the mathematical model that generated the data. Nonetheless, when training PINNs, choosing the correct hyperparameters is often the result of a trial and error procedure; when dealing with high dimensional problems, hyperparameter tuning can become unfeasible. To date, the only application of PINNs to 4D flow MRI data is the work of Fathi et al. [9]. The authors tested their approach on synthetic 4D flow data from a reference computational fluid dynamics (CFD) simulation for which ground truth velocity and pressure fields were known, achieving SR and superior denoising with respect to DFWs [21]. Successively, they evaluated their method on *in vitro* 4D flow velocity measurements, but not on *in vivo* medical data.

PINNs can be thought of as a subset of coordinate-based MLPs [34,36,19]. These networks can represent complex signals by taking low-dimensional coordinates as input and returning the value of the signal at the input locations. Such signal representations are often referred to as INRs [34]. Signals represented by coordinate-based MLPs can be orders of magnitude more compact than their grid-based counterparts [37]. However, coordinate-based MLPs are known to suffer from severe spectral bias, namely they struggle to learn the high frequency signal components. Successful biomedical applications of coordinate-based networks include 3D vascular surface reconstruction [1] and non-rigid medical image registration [44]. A further improvement of coordinate-based MLPs to capture high order derivatives in the output signal was achieved by employing sinusoidal activation functions for every hidden neural network layer, as introduced by Sitzmann et al. [34]. Sinusoidal representation networks (SIRENS) have been shown to be suited for representing complex natural signals, including images, solutions to Poisson equations and 3D shapes [34].

In this work, we employ SIRENS for learning time-varying velocity fields measured by 4D flow MRI. We leverage the MLP's architecture to introduce an implicit prior to constrain the space of solutions and investigate such implicit regularization bias towards lower frequencies, which simultaneously prevents overfitting and reduces noise in flow-encoded MR images.

## 2. Method


### 2.1. Problem setting

Let  $V^* \in \mathbb{R}^{N_r \times N_c \times N_s \times N_t \times 3}$  be a time-resolved flow image volume sequence obtained from a reconstructed 4D flow acquisition, with  $N_r$ ,  $N_c$ ,  $N_s$  and  $N_t$  being the number of rows, columns, slices and cardiac frames, respectively. Let  $\Delta x \in \mathbb{R}^3$  and  $\Delta t \in \mathbb{R}$  be the physical spacings between adjacent voxels along the spatial and temporal dimensions. At a voxel center with 4D coordinates  $[x_i, t_j]$ , where  $x_i \in \mathbb{R}^3$ , we denote the measured velocity vector as:  $V^*[x_i, t_j] \in \mathbb{R}^3$ , with  $i \in [1, N_r \times N_c \times N_s]$ , and  $j \in [1, N_t]$ . Herein, the index  $i$  refers to the flattened spatial coordinates of the image (i.e., converted into a one-dimensional array, merging all elements into a single level).

We are interested in representing  $V^*$  with a continuous function  $f: \mathbb{R}^4 \rightarrow \mathbb{R}^3$ , where  $f(x_i, t_j) = V^*[x_i, t_j]$ . To approximate  $f$ , we use an MLP  $f_{\Theta}$  with weights  $\Theta$  and with sinusoidal activation functions (SIREN). The  $l$ -th layer of a SIREN receiving a generic input tensor  $\mathbf{x}_l \in \mathbb{R}^{Q_l}$  performs the following operation:

$$\mathbf{x}_{l+1} = \sin(\Theta_l \mathbf{x}_l + \mathbf{b}_l), \quad (4)$$

where  $\Theta_l \in \mathbb{R}^{P_l \times Q_l}$  and  $\mathbf{b}_l \in \mathbb{R}^{P_l}$  are the weight matrix and biases of the  $l$ -th layer, respectively. Following [34], each weight  $\theta$  is initialized so that  $\theta \sim \mathcal{U}(-\sqrt{6}/c, \sqrt{6}/c)$ , where  $c$  is the generic input feature size. As motivated by [34], this initialization makes the input of each sine activation being Gauss-Normal distributed, and the output approximately arcsine-distributed. Furthermore, as proposed by [34], the first layer



**Fig. 1.** Conceptual diagram of the proposed approach. The SIREN model is given spatial coordinates  $\hat{x}_i, \hat{y}_i, \hat{z}_i$  and temporal coordinate  $\hat{t}_j$ . The network outputs the three velocity components evaluated at the input coordinates denoted as  $u_{ij}, v_{ij}, w_{ij}$ .

of the SIREN is modified as:  $\sin(\omega_0 \cdot \Theta \mathbf{x} + \mathbf{b})$ . Following [34], we set  $\omega_0 = 30$ , ensuring that the first sine layer spans multiple periods over the input coordinates interval.

For our purposes, a key advantage of this formulation lies in representing a high dimensional image as a continuous function that can be queried at arbitrary spatio-temporal resolutions.

## 2.2. Training a SIREN in 4D

In most practical cases dealing with blood vessels, one is only interested in reconstructing blood flow within a bounded region  $\Omega \subset \mathbb{R}^3$  with inflow boundary  $\Gamma_i$ , outflow boundary  $\Gamma_o$  and wall boundary  $\Gamma_w$ , and within a time interval  $[t_a, t_b]$ . The inner fluid region is denoted by  $\Omega_f = \Omega \setminus \{\Gamma_i \cup \Gamma_o \cup \Gamma_w\}$ . In our approach, velocity field reconstruction is achieved by sampling  $N_f$  spatial voxel coordinates  $\mathbf{x}^{(f)}$  from  $\Omega_f$  repeated over the time interval  $[t_a, t_b]$ , i.e.,  $(\mathbf{x}_i^{(f)}, t_j)$ , with  $i = 1, 2, \dots, N_f; j = 1, 2, \dots, N_t$ . Additionally, we enforce the no-slip condition on the vessel wall by sampling  $N_w$  spatial coordinates from  $\Gamma_w$  repeated over the time interval  $[t_a, t_b]$ , i.e.,  $(\mathbf{x}_p^{(w)}, t_q)$ , with  $p = 1, 2, \dots, N_w; q = 1, 2, \dots, N_t$ . Of note, we oversample spatial coordinates from  $\Gamma_w$  by setting  $N_f \approx N_w$  sample  $N_w \approx N_f$  spatial coordinates from  $\Gamma_w$ . We denote the total number of spatial coordinates used for SIREN training as  $N = N_f + N_w$ . Unlike the original SIREN formulation, this approach allows us to only use a relatively small set of coordinates compared to the total number of image points, greatly reducing training time and memory cost.

For each generic input coordinate pair  $(\mathbf{x}, t)$ , the following non-dimensionalization is performed:

$$\hat{\mathbf{x}} = \frac{\mathbf{x} - \mathbf{x}_{min}}{\Delta \mathbf{x}} D, \quad \hat{t} = \frac{t - t_{min}}{\Delta t} D, \quad (5)$$

where  $\mathbf{x}_{min}$  and  $t_{min}$  are the minimum spatial and temporal coordinates, and we set  $D = 0.01$  thus ensuring that the first sine layer of our MLP spans multiple cycles in the input coordinate interval.

Training  $f_\Theta$  implies solving the following minimization problem:

$$\min_{\Theta} \mathcal{L}(\Theta), \quad (6)$$

where the loss function  $\mathcal{L}$  is given by the misfit between the MLP prediction and the measured data plus a boundary condition term:

$$\mathcal{L}(\Theta) = \sum_{i=1, j=1}^{N_f, N_t} \|f_\Theta(\hat{\mathbf{x}}_i^{(f)}, \hat{t}_j) - V^*[\mathbf{x}_i^{(f)}, t_j]\|_2^2 + \sum_{p=1, q=1}^{N_w, N_t} \|f_\Theta(\hat{\mathbf{x}}_p^{(w)}, \hat{t}_q)\|_2^2. \quad (7)$$

Hence, the only supervision comes from the image values and fixed Dirichlet boundary conditions represented by the no-slip condition at

the wall, making the approach fully unsupervised. A diagram representation of the proposed method is shown in Fig. 1. Training is carried out using a limited memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) algorithm with learning rate of 1 until  $\nabla_{\Theta} \mathcal{L} = 0$ .

## 2.3. SIREN evaluation

Once trained,  $f_\Theta$  can be queried at arbitrary spatio-temporal collocation points denoted as  $(\mathbf{x}', t')$ , sampled from the spatio-temporal domain  $\Omega \times [t_a, t_b]$ . To evaluate the generalization capabilities of a trained SIREN, collocation points corresponding to  $N'$  spatial coordinates and  $N'_t$  temporal coordinates, i.e.,  $(\mathbf{x}'_n, t'_n)$ , with  $n = 1, 2, \dots, N'; m = 1, 2, \dots, M'$ , are defined at  $\approx \times 20$  higher spatial resolution and  $\times 10$  higher temporal resolution than image coordinates. Non-dimensionalization of evaluation coordinates is operated consistently with the one used for the training set:

$$\hat{\mathbf{x}}' = \frac{\mathbf{x}' - \mathbf{x}_{min}}{\Delta \mathbf{x}} D, \quad \hat{t}' = \frac{t' - t_{min}}{\Delta t} D. \quad (8)$$

Therefore, we evaluate  $f_\Theta$  on:  $(\hat{\mathbf{x}}'_n, \hat{t}'_m), n = 1, \dots, N'; m = 1, \dots, M'$ .

## 2.4. Error quantification

To quantify errors obtained in experiments, three different metrics were used. Differences between a reference vector field  $\mathbf{u}_{ref}$  and another generic vector field  $\mathbf{u}$ , were evaluated by computing velocity normalized root mean square error (vNRMSE), magnitude normalized root mean square error (mNRMSE), and the direction error (DE) as:

$$mNRMSE = \frac{1}{\max |\mathbf{u}_{ref}|} \sqrt{\frac{1}{K} \sum_{k=1}^K (|\mathbf{u}| - |\mathbf{u}_{ref}|)_k^2 \times 1000}, \quad (9)$$

$$vNRMSE = \frac{1}{\max |\mathbf{u}_{ref}|} \sqrt{\frac{1}{K} \sum_{k=1}^K (\mathbf{u} - \mathbf{u}_{ref})_k^2 \times 1000}, \quad (10)$$

$$DE = \frac{1}{K} \sum_{k=1}^K \left( 1 - \frac{|\mathbf{u}_{ref, k} \cdot \mathbf{u}_k|}{|\mathbf{u}_{ref, k}| |\mathbf{u}_k|} \right) \times 100, \quad (11)$$

where  $K$  is the generic number of 4D points where the two velocity fields are evaluated.

## 2.5. Wall shear stress analysis

From a generic velocity field, the WSS field was calculated following the approach described in [22]. From the definition of WSS for a Newtonian fluid:




Fig. 2. a) Computational domain. b) Time-varying 3-directional velocity profiles prescribed at as inlet boundary condition. c) Flow waveform imposed at  $\Gamma_i$ .

$$WSS = \mu \left( \frac{\partial v}{\partial y} \right)_{y=0}, \quad (12)$$

where  $\mu$  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, for each spatial 3D point on the vessel wall ( $\Gamma_w$ ), the implemented WSS calculation method requires interpolation of the velocity fields at 2 points evenly spaced by a distance  $\delta_n$  along the inward normal. The so obtained local velocity profile is interpolated with a quadratic function, whose analytical derivative is used to approximate  $\partial v / \partial y$ . In our experiments, we set  $\delta_n = 0.5 \text{ mm}$  and the no-slip condition at the wall.

## 2.6. Case 1: synthetic 4D flow MRI

**CFD simulation** To have a benchmark for evaluating the proposed method, a synthetic 4D flow MRI acquisition was created from a reference CFD simulation. First, the ascending aorta of a subject with thoracic aortic aneurysms (TAA) was segmented from 3D magnetic resonance angiography (MRA) images using open-source software [45]. The segmented domain  $\Omega$ , was divided into 3 subdomains: inlet  $\Gamma_i$ , outlet  $\Gamma_o$  and wall  $\Gamma_w$  (Fig. 2a). A 3D tetrahedral mesh with a base size of 0.6 mm and minimum size of 0.05 mm was generated using the vascular modeling toolkit (VMTK) library [15]. The final volumetric mesh consisted of  $\approx 800\text{k}$  nodes and 4M elements. Time-varying 3-directional velocity profiles (Fig. 2b) were prescribed as inlet boundary conditions, mapping a realistic TAA inlet velocity profile to the  $\Gamma_i$  following the approach described in [30] and producing the flow waveform represented in Fig. 2c. A zero-pressure condition was enforced on  $\Gamma_o$  and a homogeneous Dirichlet boundary condition (no-slip) was assumed on  $\Gamma_w$ . Blood was modeled as a Newtonian fluid with constant density  $\rho = 1060 \text{ kg/m}^3$  and dynamic viscosity  $\mu = 0.004 \text{ Pa}\cdot\text{s}$ . A finite volume simulation was run at a fixed time step of 0.001 s. An implicit scheme with splitting of operators (PISO) was used to solve the governing equations of blood flow in Star-CCM+. To lighten the computational cost related to graphics processing unit (GPU) memory, and to focus on the more interesting flow details, results were exported at every timestep within the time interval [0.2–0.399] s (peak to late systole), yielding a sequence of  $M' = 200$  velocity fields  $\mathbf{u}^*$  defined on the computational nodes  $\mathbf{x}'_n, n = [1, \dots, N']$  over the simulation time coordinates  $t'_m, m = [1, \dots, M']$ .

**Synthetic data creation** Noise-free, high resolution CFD velocity fields were processed to obtain low resolution velocity fields corrupted with noise typical of 4D flow MRI measurements. To achieve this, the following steps were implemented partially following [10]:

1. the sequence of CFD solution snapshots denoted by  $\mathbf{u}^*$  was temporally downsampled to a sequence of  $M = \frac{M'}{h}$  velocity fields  $\bar{\mathbf{u}}$ ,

using a moving average such that:  $\bar{\mathbf{u}}(\mathbf{x}'_n, t'_j) = \frac{1}{h} \sum_{k=1}^{m+h-1} \mathbf{u}^*(\mathbf{x}'_n, t'_k)$ , with  $h$  being the time averaging window size;

2. each time frame of  $\bar{\mathbf{u}}$  was converted to a uniform Cartesian grid with voxel size of  $1 \text{ mm}^3$  using a linear interpolation scheme to assign velocity vector values to grid cells, yielding a sequence of Cartesian grids  $\tilde{U}[\tilde{\mathbf{x}}_i, t_j] \in \mathbb{R}^{N_r \times N_c \times N_s \times N_t \times 3}$ , with  $N_r = 76$ ,  $N_c = 112$  and  $N_s = 292$ ;
3. each velocity grid in  $\tilde{U}$  was converted to a complex tensor containing magnitude and phase images using suitable VENC values;
4. the fast Fourier transform was applied to obtain the corresponding  $k$ -space data;
5. a truncation of the 3D  $k$ -space data (high frequencies) was performed, effectively decreasing the spatial resolution by a factor of 2;
6. a zero-mean Gaussian noise with standard deviation  $\sigma$  corresponding to the desired SNR (calculated according to [10]) was added to the  $k$ -space data;
7. a randomized sampling mask drawn from a normal distribution and covering 5% percent of the  $k$ -space was applied to further under-sample the frequency content, keeping a fully sampled calibration region of  $5 \times 5 \times 5$  in the center of  $k$ -space;
8. the inverse Fourier transform was applied to the undersampled, noise-corrupted  $k$ -space, yielding a complex tensor of magnitude and phase images;
9. complex images were converted back to real images of velocity fields using VENC values consisted with step 3, obtaining a sequence of noisy synthetic velocity measurements  $U^*$  defined at voxel coordinates  $\mathbf{x}_i, i = 1, \dots, N$  with isotropic voxel size of  $2 \times 2 \times 2 \text{ mm}^3$  and at times  $t_j, j = 1, \dots, M$ .

For each velocity direction, the VENC value was chosen 10% larger than the maximum velocity, so that the phase wrapping/unwrapping would not introduce aliasing artifacts corresponding to a value of 120 cm/s in each direction. The level of degradation of our synthetic measurements is regulated by three parameters:  $h$ , SNR and  $S$ . We chose  $h = 40$  to obtain  $\Delta t = 0.04 \text{ s}$ , which is compatible with real medical measurements. We tested our approach on three different levels of degradation, *mild*, *medium* or *extreme* reported in Table 1. For all noise levels, velocity measurements were characterized by an isotropic voxel spacing of 2 mm for each dimension. For the *extreme* noise level, the generated synthetic images for the three different velocity components are shown in Fig. 3.

As mentioned in Sections 2.2 and 2.3, the training and test coordinates correspond to image voxel centers and to the CFD mesh nodes, respectively. In particular, the total number of spatial coordinates used for training was  $N \approx 40000$ , while the trained SIREN models were evaluated at  $N' \approx 800000$  spatial coordinates. On the other hand, the number






Fig. 3. Ground truth velocity vector components from CFD (first 3 rows), and corresponding velocity components after *extreme* noise addition. LR: left-to-right, AP: anterior-to-posterior, FH: foot-to-head.

Table 1

Degradation parameters for generating synthetic measurements from the reference velocity field.

Noise level	SNR	$S$
<i>mild</i>	20	99
<i>medium</i>	5	95
<i>extreme</i>	2	68

of temporal coordinates used for training was  $N_t = 5$ , while the SR performance was evaluated at  $M' = 50$  time points.

### 2.6.1. Comparison with existing methods

The denoising and SR performances of our approach were compared against existing methods. Among them, only 4D RBFs lend themselves to the joint task of denoising and SR. Hence, pure denoising approaches were combined with simple interpolation schemes to achieve SR. The following methods were tested:

1. linear interpolation (LITP) as a baseline for SR;
2. DFW with automatic threshold selection based on *SureShrink* [21] for denoising and LITP for SR. Velocity fields generated by this approach will be denoted as DFW;
3. 3D divergence-free radial basis functions (DF-RBFs) as described in [8] for denoising and LITP for SR. Velocity fields generated by this approach will be denoted as DF-RBF;
4. an approach based on 4D RBFs [38], but for which no official implementation was available. Hence, we implemented our version of 4D RBFs with multi-quadric kernel and local support for denoising and SR. We set 10 as the number of nearest neighboring points for local kernel support. A *soft* enforcement of the no-slip condition on the vessel wall was applied by setting null velocity values

at point belonging to the wall region. Velocity fields generated by this approach will be denoted as 4D-RBF.

### 2.7. Case 2: in vivo 4D flow MRI

A thoracic 4D flow MRI scan of a subject with ascending thoracic aortic aneurysm was retrospectively retrieved. Images were fully deidentified and provided by Weill Cornell Medicine, (NY, USA). A respiratory compensated technique was adopted with the following settings: spatial resolution (voxel size)  $1.14 \text{ mm} \times 1.14 \text{ mm} \times 0.9 \text{ mm}$ , field of view = 360 mm, flip angle =  $15^\circ$ , VENC = 200 cm/s in all 3 directions, time between consecutive frames = 30 ms, for a total of 20 frames per cardiac cycle. DICOM images were processed using open-source code [30] to compute the phase-contrast magnetic resonance angiography (PC-MRA) image and extract the segmentation of the enlarged ascending aortic tract (Fig. 4). The segmentation was performed on a time-averaged configuration by an experienced operator using a combination of a region growing semi-automatic algorithm and manual delineation of the vessel contours. As for Case 1, a volumetric mesh with base element size of 0.6 mm and approximately 4M elements was generated. Mesh node coordinates were used for evaluating the trained model, entailing an  $\approx \times 20$  increase between training and evaluation spatial coordinates. As for Case 1, temporal coordinates were increased  $\times 10$  from training to evaluation.

The velocity components measured through flow MRI can be visualized in Fig. 5.

## 3. Results

### 3.1. Case 1: hyperparameter tuning

The effects of SIREN's depth (number of layers) and width (number of neurons per layer) on the denoising and SR performances were as-




Fig. 4. a): segmented 3D geometry (red) superimposed on a sagittal slice of PC-MRA images. b) the segmented 3D region with an axially oriented slice (top) and the contours of the static segmentation superimposed on the velocity magnitude images at different time points denoted in the top left corners.




Fig. 5. Sagittal 2D slice within the aortic aneurysm (bottom left), together with colormaps of left-to-right (LR rows), posterior-to-anterior (PA rows) and foot-to-head (FH rows) velocity components. Rows 4 to 6 show velocities sampled on a 3D mesh by linear interpolation. Columns from left to right correspond to increasing time points.

**Table 2**

Velocity fields comparison at training coordinates. In each cell values of mNRMSE (top), vNRMSE (middle) and DE (bottom) are reported. The first column contains the errors computed between the ground truth CFD solution and the noisy 4D flow data at voxel centers. The second column contains errors computed between the ground truth CFD solution and the evaluation of the trained SIREN at voxel centers. Errors are reported for the three levels of noise.

	$u_{ref} = u^*(x, t)$ $u = V^*[x, t]$	$u_{ref} = u^*(x, t)$ $u = f_{\Theta}(x, t)$
<i>mild</i>	31.7	30.6
	28.0	27.1
	27.2	27.0
<i>medium</i>	53.2	38.3
	52.4	39.4
	33.0	31.1
<i>extreme</i>	70.9	45.1
	82.8	50.7
	37.9	32.7

**Table 3**

Effect of SIREN number of layers (depth) and number of neurons per layer (width) on denoising and super-resolution using data with *mild* noise level. In each cell values of mNRMSE (top), vNRMSE (middle) and DE (bottom) are reported.

		Depth				
		4	8	12	16	20
Width	100	3.52	3.56	3.53	3.42	3.46
		3.43	3.36	3.32	3.22	3.27
		6.8	6.4	6.3	6.22	6.31
	200	3.48	3.46	3.52	3.52	<b>3.41</b>
		3.36	3.27	3.27	3.27	<b>3.21</b>
		6.43	6.23	6.16	6.16	6.14
	300	3.51	3.53	3.62	3.50	3.43
		3.40	3.30	3.33	3.24	3.21
		6.47	6.24	6.19	6.14	6.13
400	3.49	3.47	3.60	3.49	3.53	
	3.34	3.25	3.31	3.25	3.29	
	6.37	6.16	6.17	<b>6.12</b>	6.15	
500	3.59	3.47	3.55	3.54	3.45	
	3.45	3.26	3.30	3.29	3.23	
	6.42	6.19	6.17	6.19	6.14	

essed by training different configurations of  $f_{\Theta}$  on synthetic images with *mild*, *medium* and *extreme* noise level. Trained models were evaluated on CFD nodal mesh coordinates  $x'$  and at time points  $t'$  evenly spaced by 0.004 s, effectively oversampling  $\Omega$  by  $\approx \times 20$  and the time interval [0.2–0.399] by  $\times 10$ . Results are reported in Tables 3, 4 and 5, for *mild*, *medium* and *extreme* noise levels, respectively. For *mild* noise levels, all the tested SIREN configurations gave low errors with respect to ground truth velocity fields. Networks with greater width and depth resulted in only slightly lower mNRSME, vNRMSE and DE. For *medium* noise levels, better results were obtained by wider models, but not necessarily by deeper ones. For these noise settings, the best performing SIREN consisted of 12 layers, each with 500 neurons. In the case of *extreme* noise, wider networks gave worse results than narrower ones, while deeper architectures generally produced more accurate velocity fields compared to ground truth. Wider models showed a tendency to overfit high frequency noise, as shown in Figs. 6.

We chose the single best configuration as the one that minimized the sum of vNRMSE, mNRSME and DE for all three noise levels, hence a SIREN 20 layer deep and with 300 neurons per layer was selected for

**Table 4**

Effect of SIREN number of layers (depth) and number of neurons per layer (width) on denoising and super-resolution using data with *medium* noise level. In each cell values of mNRMSE (top), vNRMSE (middle) and DE (bottom) are reported.

		Depth				
		4	8	12	16	20
Width	100	4.48	4.63	4.65	5.86	5.29
		4.73	4.92	4.94	6.13	5.61
		10.5	10.8	10.9	13.6	12.3
	200	3.98	3.98	4.42	4.49	4.85
		4.24	4.24	4.72	4.81	5.15
		9.66	9.61	10.5	10.7	11.5
	300	3.94	3.94	3.97	4.02	4.20
		4.20	4.20	4.21	4.30	4.51
		9.59	9.56	9.5	9.67	10.21
400	4.03	<b>3.81</b>	3.89	3.87	4.03	
	4.35	4.09	4.16	4.12	4.28	
	9.87	9.39	9.48	9.44	9.69	
500	3.95	3.82	3.83	3.85	3.84	
	4.24	4.11	<b>4.08</b>	4.10	4.10	
	9.62	9.41	9.38	<b>9.33</b>	9.36	

**Table 5**

Effect of SIREN number of layers (depth) and number of neurons per layer (width) on denoising and super-resolution using data with *extreme* noise level. In each cell values of mNRMSE (top), vNRMSE (middle) and DE (bottom) are reported.

		Depth				
		4	8	12	16	20
Width	100	5.31	5.37	5.26	5.30	5.87
		5.73	5.80	5.73	5.75	6.29
		13.01	13.33	12.93	13.02	14.17
	200	5.84	5.99	5.16	<b>5.15</b>	5.21
		6.62	6.91	5.68	5.67	5.64
		14.83	15.11	12.87	12.78	12.84
	300	7.30	7.74	5.40	5.21	5.16
		8.50	9.05	6.09	5.77	<b>5.63</b>
		18.86	19.91	13.41	12.95	<b>12.7</b>
400	8.86	8.14	7.61	5.56	5.44	
	10.3	9.52	8.89	6.33	6.10	
	22.34	21.02	19.68	13.98	13.51	
500	9.41	8.07	7.95	7.46	6.50	
	10.9	9.45	9.28	8.65	7.56	
	23.52	20.88	20.5	19.16	16.66	

all downstream comparisons against other methods and on the medical dataset.

The trained model with the chosen architecture was also evaluated at the training coordinates  $(x, t)$  where noisy velocity observations were fitted, obtaining an increase in accuracy compared to the data values themselves. Quantitative results are reported in Table 2.

### 3.2. Case 1: comparison with existing methods

#### 3.2.1. Velocity fields

The SIREN configuration that showed the best results was evaluated on the synthetic images with three different levels of noise. Velocity fields obtained with our approach were compared against the existing methods listed in Section 2.6.1. The proposed SIREN gave lower mNRSME, vNRMSE and DE for all noise levels (Table 6). Among the LITP-based methods, DF-RBF [8] performed better than LITP alone,



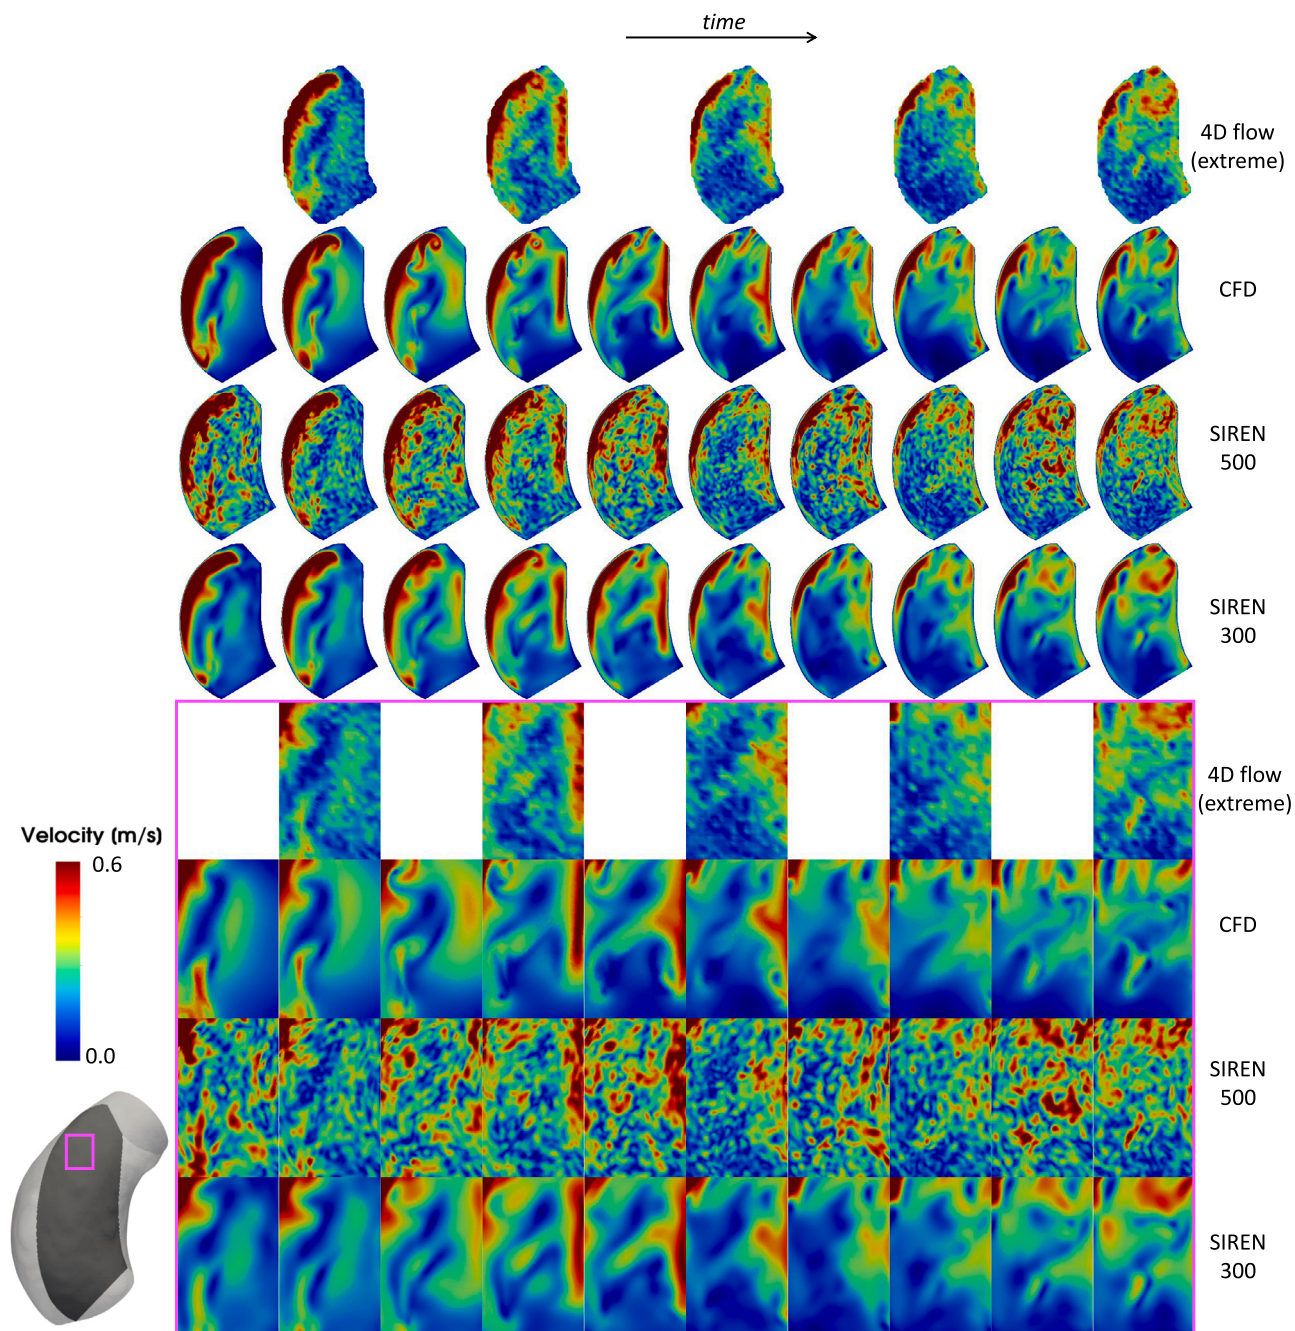


Fig. 6. Sagittal 2D slice within the ascending aorta (bottom left), together with a zoomed in region shown in pink. Rows 1 and 5: images with *extreme* noise added; rows 2 and 6: CFD solution; rows 3 and 7: results of a SIREN with 20 layers and 500 neurons per layer; rows 4 and 8: results of a SIREN with 20 layers and 300 neurons per layer. Rows 5 to 8 represent the zoomed in region in contoured in pink.

while DFW provided the worst performance. Qualitatively, all existing methods produced very similar velocity fields for the *mild* noise level case, showing suboptimal capabilities of reconstructing finer flow details (Fig. 7). For the same level of noise, our method was able to fit flow details more accurately. Similar results were observed for the *medium* and *extreme* noise levels. As shown in Figs. 8 and 9, all baseline approaches tended to perform denoising by oversmoothing the data. On the other hand, our method was able to filter out noise but better preserving finer flow structures.

### 3.2.2. Wall shear stress fields

The same methodology described in Section 2.5 was applied to the velocity fields obtained by CFD (ground truth), DFW, DF-RBF, 4D-RBF

and our best SIREN (20 layers, 300 neurons per layer). Results obtained by the different approach are reported in Table 7. Our method gave the lowest mNRSME, vNRMSE and DE, outperforming the others for all the tested noise levels. WSS fields computed from the denoised and super-resolved velocity field obtained by the different approaches can be visualized as 3D colormaps in Figs. 10, 11 and 12 for *mild*, *medium* and *extreme* noise levels, respectively. More visible differences produced by the different methods can be appreciated for the *extreme* noise level (Figs. 12). In this case, all methods, including ours, underestimated WSS magnitudes. Our SIREN showed superior performance, giving WSS fields in good agreement with CFD-derived results, with high WSS magnitudes on the outer curve of the vessel wall, as expected.






Fig. 7. Velocity colormaps on a 2D sagittal slice (bottom left). Top row: synthetic measurements with *mild* noise level at time points evenly spaced by  $\Delta t = 40$  ms. Second row: ground truth velocity fields from CFD at time points evenly spaced by  $\Delta t = 20$  ms. Third row: LITP results. Fourth row: DFW results. Fifth row: DF-RBF results. Sixth row: 4D-RBF results. Seventh row: our method.

**Table 6**  
Velocity field comparison with existing methods. In each cell values of mNRMSE (top), vNRMSE (middle) and DE (bottom) are reported.

	LITP	DFW + LITP	DF-RBF + LITP	4DRBF	SIREN
<i>mild</i>	5.88	6.47	5.90	6.91	<b>3.50</b>
	4.96	5.40	4.99	5.68	<b>3.24</b>
	7.13	7.96	7.2	8.15	<b>6.14</b>
<i>medium</i>	6.11	7.26	6.05	6.79	<b>4.02</b>
	5.50	6.18	5.42	6.01	<b>4.30</b>
	10.1	10.16	9.78	10.53	<b>9.67</b>
<i>extreme</i>	7.06	8.53	6.85	7.57	<b>5.21</b>
	6.96	7.43	6.62	7.04	<b>5.77</b>
	14.16	13.01	13.26	13.9	<b>12.95</b>

**Table 7**  
WSS field comparison with existing methods. In each cell values of mNRMSE (top), vNRMSE (middle) and DE (bottom) are reported.

	LITP	DFW + LITP	DF-RBF + LITP	4DRBF	SIREN
<i>mild</i>	8.76	8.65	8.68	18.6	<b>5.62</b>
	6.03	6.10	6.00	12.31	<b>4.02</b>
	3.43	4.64	3.52	3.84	<b>2.1</b>
<i>medium</i>	8.43	7.93	8.29	18.6	<b>6.53</b>
	6.05	5.99	5.96	12.6	<b>4.76</b>
	4.94	5.83	4.86	6.85	<b>3.98</b>
<i>extreme</i>	10.32	9.06	10.1	19.5	<b>7.82</b>
	7.59	7.01	7.41	13.4	<b>6.04</b>
	7.74	8.41	7.5	10.34	<b>7.0</b>




Fig. 8. Velocity colormaps on a 2D sagittal slice (bottom left). Top row: synthetic measurements with *medium* noise level at time points evenly spaced by  $\Delta t = 40$  ms. Second row: ground truth velocity fields from CFD at time points evenly spaced by  $\Delta t = 20$  ms. Third row: LITP results. Fourth row: DFW results. Fifth row: DF-RBF results. Sixth row: 4D-RBF results. Seventh row: our method.

### 3.3. Case 2: velocity field assessment

Table 8 shows quantitative measurements from raw MRI data and reconstructed velocity fields with the proposed SIREN with 20 layers and 300 neurons per layer. Overall, our method gave high resolution velocity fields that maintained low discrepancies in macroscopic quantitative measurements that are considered to be reasonably accurate when assessed from unprocessed MRI measurements [13], giving differences in mean and maximum flow rate of  $< 5\%$  and underestimating reverse flow index (RFI) by 9.4%, as computed in [29]. Qualitatively, our approach produced cleaner velocity fields, but maintaining the high velocity regions observed at the extrados of the ascending aorta as in the measured data (Figs. 13 and 14). From the 3D vector visualization (third and fourth rows in Figs. 13 and 14), it can be appreciated how the reconstructed velocity vector field kept the swirling patterns formed in the middle ascending aorta, but filtered out the spurious vector com-


Table 8

Quantitative measurements for Case 2 on raw flow MR measurements and on velocity fields reconstructed with our method.

	Measured data	SIREN	$\Delta\%$
Mean flow rate [L/min]	4.66	4.88	4.9%
Max flow rate [L/min]	27.7	27.2	1.8%
RFI [%]	21	19	9.4%

ponents in the near-wall regions. Similarly, the DF-RBF and 4D-RBF methods (Fig. 15, rows 4 and 5) gave visually denoised velocity fields, free from the high-frequency noise features that were observed in the unprocessed data, LITP and DFW results (Fig. 15, rows 1, 2 and 3).





**Fig. 9.** Velocity colormaps on a 2D sagittal slice (bottom left). Top row: synthetic measurements with *extreme* noise level at time points evenly spaced by  $\Delta t = 40$  ms. Second row: ground truth velocity fields from CFD at time points evenly spaced by  $\Delta t = 20$  ms. Third row: LITP results. Fourth row: DFW results. Fifth row: DF-RBF results. Sixth row: 4D-RBF results. Seventh row: our method.

### 3.4. Case 2: wall shear stress field assessment

WSS fields reconstructed by the SIREN applied on the clinical measurements were consistent with observations made for the synthetic case. Specifically, the outer curve of the aneurysm exhibited a notable high-velocity region near the wall (Fig. 16). As a consequence of this accelerated flow pattern, there was a pronounced elevation in the WSS within this particular area. This phenomenon indicates a direct correlation between the high-velocity region and the resultant increased WSS, emphasizing the influence of flow dynamics on the local hemodynamic conditions within the aneurysm. Such insights underline the significance of considering the spatial relationship between flow characteristics and WSS distribution, particularly in regions characterized by abnormal vascular geometries like aneurysms, thereby offering valuable implications for understanding and potentially treating such conditions. A maximum WSS value of 20 Pa and mean time-averaged WSS (TAWSS) of 7.5 mPa were found.

### 4. Discussion

The use of 4D flow MRI in analyzing blood flow in major vessels has been widely studied, but the limitations of this imaging technique, such as noise and low spatio-temporal resolution have hindered its use in more advance velocity-based hemodynamic analysis. For instance, some studies have attempted to quantify WSS directly from flow-encoded MR images [35,23], evaluating their accuracy at different levels of image noise or spatial resolution. Although these approaches enable fast assessment of near-wall quantities, they often tend to underestimate the true WSS values with respect to numerical simulation. In pursuit of more accurate quantification of WSS fields, researchers have turned to increasingly complex numerical simulations of blood flow [11,20]. Although CFD studies are a powerful tool for estimating WSS, they require significant computational resources and can take several hours, or even days, to run due to their computational cost.

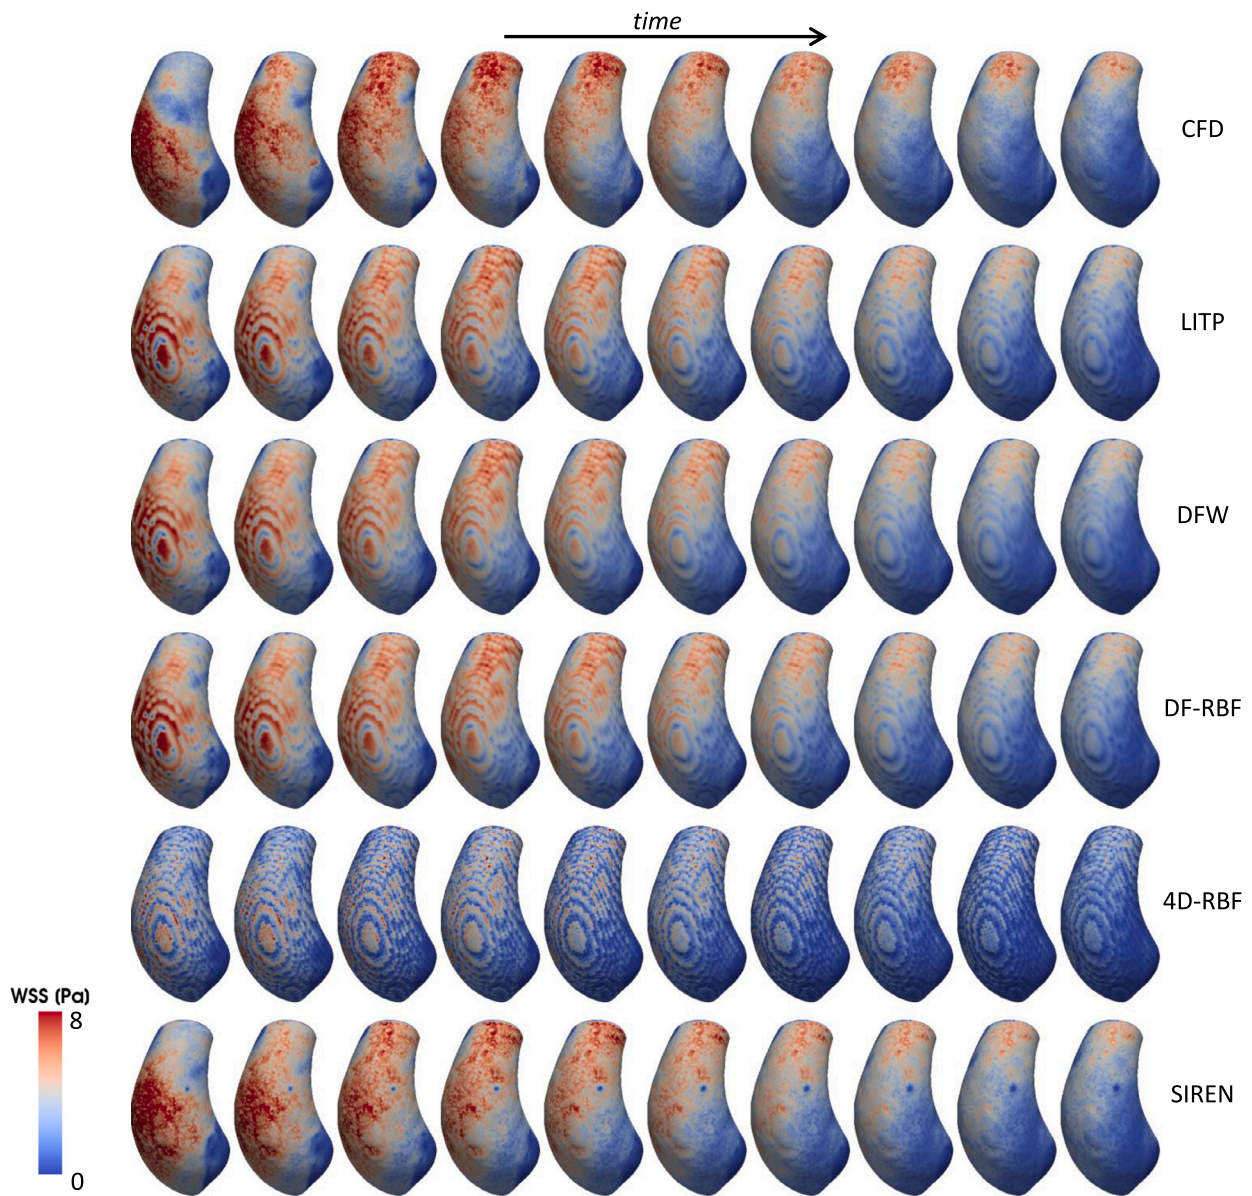


Fig. 10. Results corresponding to *mild* noise levels. WSS colormaps on the aortic wall at time points evenly spaced by  $\Delta t = 20$  ms. Top row: WSS computed on ground truth velocity fields from CFD. LITP (second row), DFW (third row), DF-RBF (fourth row), 4D-RBF (fifth row), our method (sixth row).

In this study, we proposed an unsupervised learning method based on INRs for denoising and SR of velocity fields measured by 4D flow MRI. We evaluated our approach on realistic synthetic data with various levels of noise and showed the superior performance of our method with respect to existing methods in terms of both denoising and spatio-temporal SR. The proposed approach was able to denoise and super-resolve 4D velocity data while maintaining the integrity of the dominant flow features.

Among neural network approaches, it is worth mentioning the studies of Ferdian et al. [10], Rutkowski et al. [28] and [33]. In these works, CFD simulations are used to create synthetic 4D flow MRI datasets, so to have measurements – ground truth pairs to train convolutional neural networks for denoising and SR. Hypothetically, if large realistic training data is generated, these approaches would be able to learn the denoising and SR tasks, and could generalize to new domains without the need of re-training. Nonetheless, by operating convolutions in the image domain, these models are usually engineered to super-resolve  $\times 2$  or  $\times 4$  in space alone, and they do not allow for precise evaluation of near-wall quantities. Additionally, these methods require full ground truth

supervision. For this reason, we did not include them in our comparison against existing methods. In contrast, an advantage of our approach is that the network learns a continuous representation of the velocity field which can be evaluated at any space-time location in an arbitrarily shaped domain.

The success of our method relies on two key properties of MLPs. First, we leverage SIREN's spectral bias [5,24] to achieve velocity field denoising. This property of dense fully-connected networks prevents them from learning high frequency functions. Our experiments on simulated 4D flow data (Case 1) allowed us to gain insights into the capabilities of SIRENS to denoise velocity observations. When exploiting the spectral bias for signal denoising, one should carefully choose model's width and depth such that high frequency noise is filtered out, while desirable fine signal details are maintained. Overall, with *mild* and *medium* noise levels, all the tested combinations of depths and widths were able to fit the data, with a tendency of larger models to give slightly lower errors (Tables 3, 4). For *extreme* noise levels, SIRENS with more than 300 neurons per layer overcame the spectral bias and overfitted high fre-







Fig. 11. Results corresponding to *medium* noise levels. WSS colormaps on the aortic wall at time points evenly spaced by  $\Delta t = 20$  ms. Top row: WSS computed on ground truth velocity fields from CFD. LITP (second row), DFW (third row), DF-RBF (fourth row), 4D-RBF (fifth row), our method (sixth row).

quency noise. Once the best architecture was identified, our approach outperformed all other tested denoising methods.

The second key strength of the proposed approach relies on the fact that the trained MLP fits the velocity vector field as a continuous function of space-time coordinates. In practice, this was achieved by building upon the work of Sitzmann et al. [34], who showed how SIRENS are better suited to fit complicated, feature-rich signals, such as natural images and solutions to simple PDEs. The present work is the first to adopt SIRENS for fitting 4D velocity measurements. By acting pointwise on 4D coordinates, a trained SIREN can be queried at continuous spatio-temporal locations, theoretically providing SR at arbitrarily fine spatial and temporal scales. Additionally, by incorporating time as an input feature, SIREN's point-wise outputs are implicitly affected by temporal neighboring point features. A limitation of such an approach is represented by the fact that velocity measurements are fitted in a point-wise manner as in [25]. Indeed, in 4D flow MRI, reconstructed images are a spatiotemporally averaged quantities, and a more sophisticated approach should account for such averaging by performing a convolution of the MRI system point spread function with the com-

plex phase contrast data [27]. Interestingly, the trained SIREN gave errors that were lower than the errors of the fitted data values themselves, implying that the method increases the accuracy at the training points (Table 2). This finding implies that the proposed method acts as a denoiser at both high and low-resolutions, hence the denoising performance can be attributed to the implicit spectral bias of the model rather than its smooth interpolation capabilities. In principle, our approach is closely related to the formulation introduced by PINNs [25], with the main difference lying in the definition of the loss function. A more sophisticated extension of our coordinate-based approach is represented by Fathi et al. [9], who achieve SR and denoising of 4D flow MRI by taking into consideration the underlying physics of the problem. Instead, we train our models only using a data fidelity term, neglecting physical priors. Despite representing a simplification with respect to [9], our work clearly shows that when dealing with coordinate-based networks and high-dimensional noisy data, the spectral bias of the model itself introduces some regularization, hence, future work on PINNs dealing with similar data will have to acknowledge this aspect, discerning between the regularization provided by the physics vs. the one provided



**Fig. 12.** Results corresponding to *extreme* noise levels. WSS colormaps on the aortic wall at time points evenly spaced by  $\Delta t = 20$  ms. Top row: WSS computed on ground truth velocity fields from CFD. LITP (second row), DFW (third row), DF-RBF (fourth row), 4D-RBF (fifth row), our method (sixth row).

by the architecture. Our approach leverages the concept introduced by the seminal work of [41], showing how a neural network can be used as a prior with excellent results in inverse problems such as denoising and SR. However, it is important to note that a limitation of the proposed approach is that it cannot correct all errors that are typically present in 4D flow MRI data. Essentially, our method performs regression of vector data in 4D, and it cannot add high frequency flow details that are lost during the MRI reconstruction process. Additionally, the capabilities of our method to correct for velocity aliasing errors are still unexplored. The introduction of N-S residuals in differential form in the loss function could potentially lead to more physically consistent velocity fields, but would significantly slow down the training process. Moreover, a supervised approach trained on ground truth CFD data could potentially overcome this limitation. In the present study, we were interested in demonstrating the feasibility of a method that could potentially be applied to real medical scenarios, where speed of execution is mandatory. For the TAA patient (Case 2), using an NVIDIA A100 graphics card, our

method took approximately 4 minutes to train and less than 2 seconds to evaluate at a fine spatio-temporal resolution.

In contrast to most data-based approaches, the devised method requires the definition of a bounded domain  $\Omega \times [t_a, t_b]$  (Section 2.2). On one hand, this choice represents a limitation of our workflow, entailing longer processing times. On the other hand, precise definition of a smooth vessel wall surface enables robust computation of WSS, a clinically important hemodynamic biomarker. Results on synthetic data showed good quantitative and qualitative agreement between predicted and reference CFD data (Table 7 and Figs. 10, 11, 12). Nonetheless, when comparing WSS fields obtained with different approaches, it is worth mentioning that other baseline methods, such as LITP, DFW, and DF-RBF suffer from voxel-grid aliases. On the other hand, 4D-RBF and the proposed SIREN eliminate this type of aliases, representing an important advantage compared to traditional interpolation methods.

An interesting result of our analysis on model's architecture revealed that, when trained on data with *extreme* noise levels, wider and deeper models overfitted noise and produced worse results. During training of



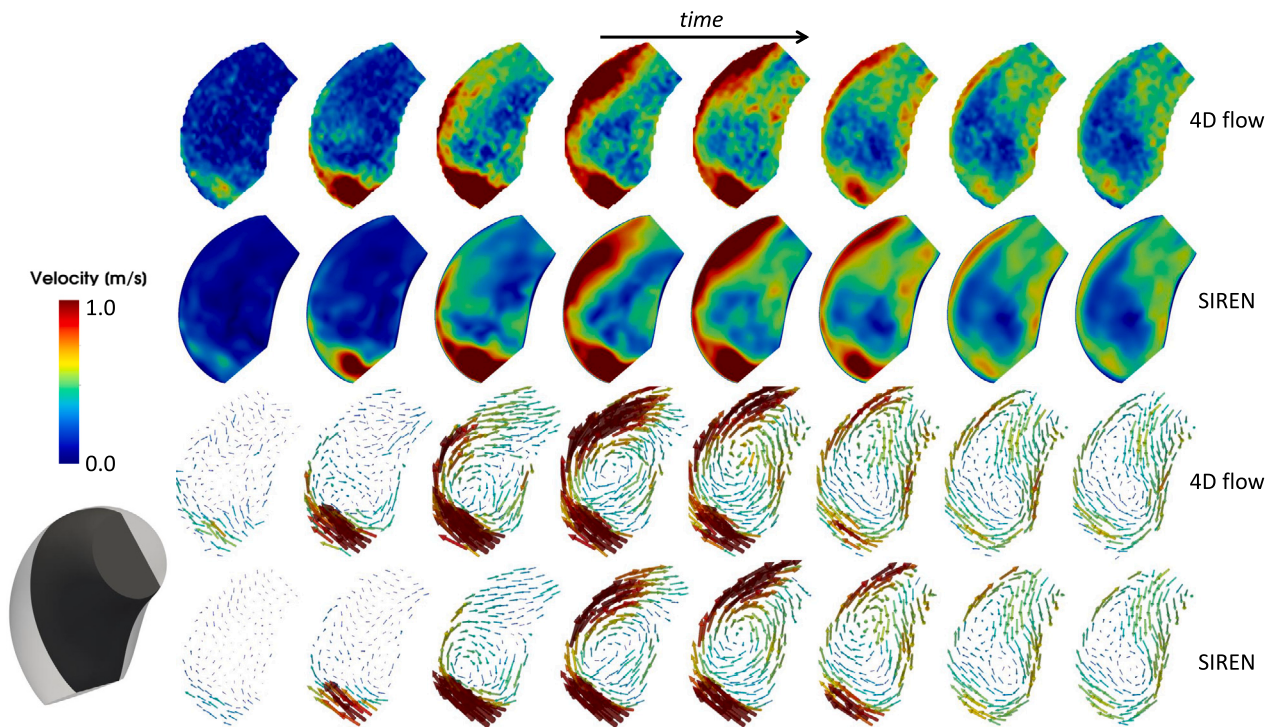


Fig. 13. Velocity magnitude colormaps (rows 1 and 2) and velocity vectors (rows 3 and 4) for Case 2 on a sagittal 2D slice (bottom left). Rows 1 and 3: 4D flow measurements, rows 2 and 4: SIREN results.

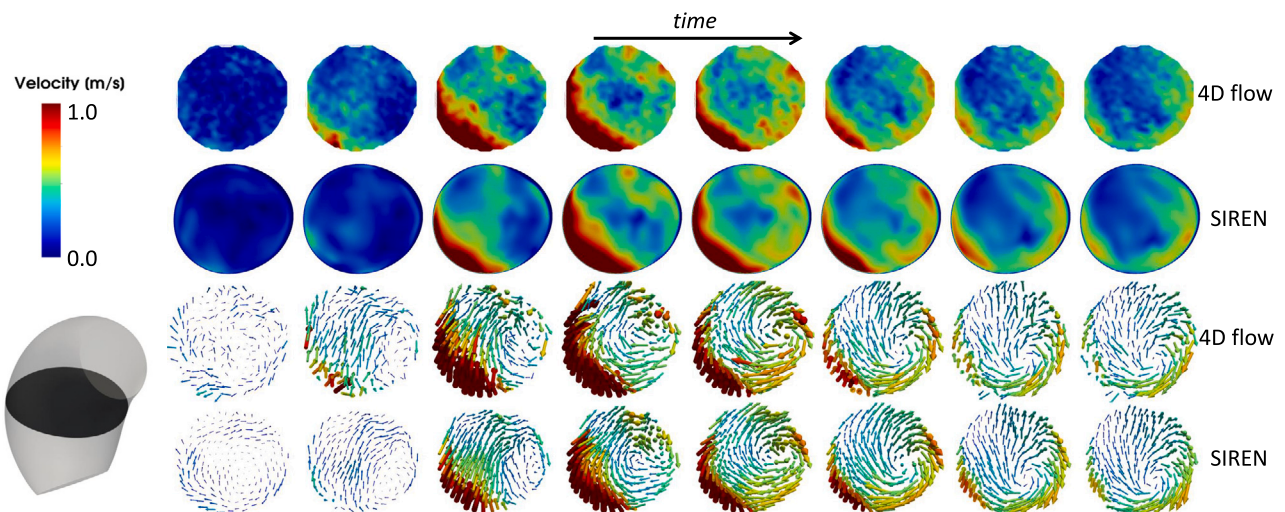


Fig. 14. Velocity magnitude colormaps (rows 1 and 2) and velocity vectors (rows 3 and 4) for Case 2 on an axially oriented 2D slice (bottom left). Rows 1 and 3: 4D flow measurements, rows 2 and 4: SIREN results.

the largest SIREN width 500, depth 20) on data with *extreme* noise levels, we observed an initial decrease and subsequent increase of the testing loss function (Fig. 17). While it is true that a lower error would have been obtained by stopping training earlier, in practice this approach is not feasible. In fact, when applying our method in real clinical settings, one cannot rely on ground truth data to compute the test loss.

Results on real medical data revealed a high WSS region on the outer curvature of the ascending aorta (Fig. 16), with most values ranging from 0 to 15 Pa. These results are in good agreement with a recent study on TAA biomechanics [32], who reported maximum WSS values of  $10.18 \pm 4.14$  Pa for a cohort of 10 patients. However, it is worth noting that any interpolation of the coarse 4D flow data is unlikely to produce exact WSS fields; the improvement in WSS accuracy

can be thought of as a consequence of the reduced noise representation compared to other existing methods. Furthermore, in our analyses, we considered a static 3D segmentation of the aortic wall. Despite a precise overlap was qualitatively observed with the velocity magnitude image (Fig. 4), for cases where wall motion is more evident, a better approach would involve a dynamic segmentation of the vessel, as this has been shown to have a significant impact on hemodynamic analysis [43]. Moreover, given the limited resolution of 4D flow MRI, determining the aortic wall position itself from this imaging modality can be challenging. In this work, the uncertainty associated with the definition of the boundary was not quantified; this aspect should be addressed by future studies focusing on WSS assessment. In principle, adopting an approach similar to [17] could allow for a joint optimization of flow re-

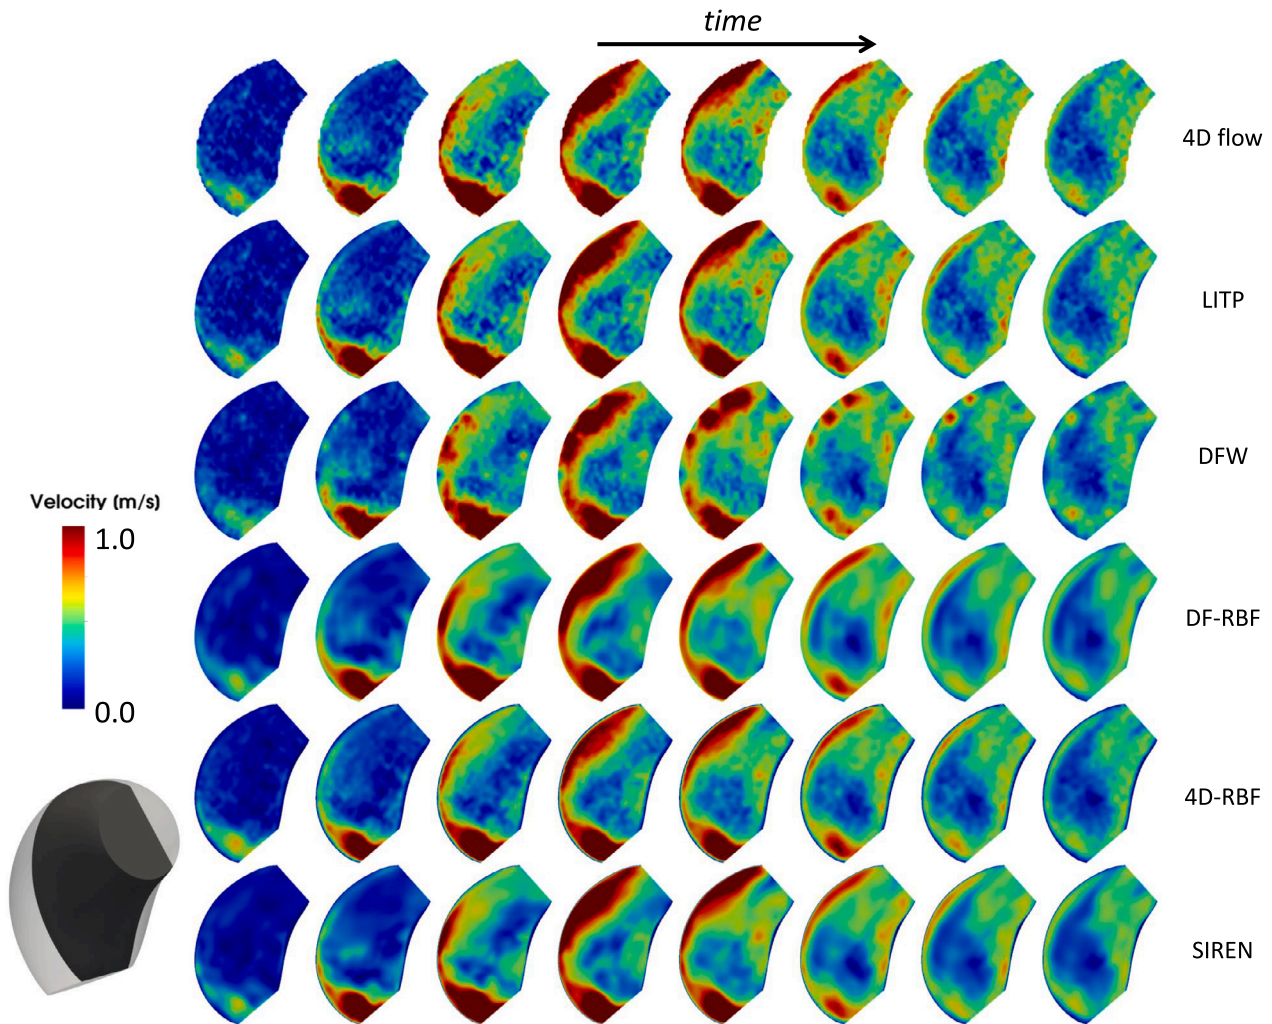


Fig. 15. Velocity magnitude colormaps on a 2D sagittal slice (bottom left). Top row: clinical 4D flow measurements. Second row: LITP results. Fourth row: DFW results. Fifth row: DF-RBF results. Sixth row: 4D-RBF results. Seventh row: our method.

construction and segmentation. Even though more detailed analyses are needed on the use of neural networks for *in vivo* WSS estimation, our results indicate that velocity fields produced by our method are suitable for extraction of derived biomarkers that are affected by noise.

Ideally, when applying our method on clinical scans, one could identify the SNR of the 4D flow images and then choose the most appropriate SIREN architecture. However, quantifying the SNR in clinical 4D flow scans is quite challenging as it is usually performed by taking the difference of two symmetrically flow-encoded magnitude images, which can be expected to have similar signal magnitude. Future research efforts could be devoted to improving the feasibility of SNR quantification from flow-encoded images, enhancing the effectiveness and applicability of the proposed approach.

#### 4.1. Conclusions

In this work we showed the feasibility of SIRENS to represent complex, high dimensional blood flow velocity fields measured by 4D flow MRI. By training on low resolution coordinates, our method is quick to execute for new cases and easy to implement. By carefully tuning our SIREN architecture, we exploit the spectral bias to obtain a functional representation of our data with reduced noise, outperforming widely adopted existing solutions. Our method provides continuous velocity fields that can be queried at arbitrary spatio-temporal locations, effectively achieving 4D super-resolution.

The code to reproduce the described methodology can be found at <https://github.com/saitta-s/INRS-4DFlowMRI>.

#### CRediT authorship contribution statement

**Simone Saitta:** Writing – review & editing, Writing – original draft, Visualization, Validation, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Marcello Carioni:** Writing – review & editing, Methodology, Formal analysis, Conceptualization. **Subhadip Mukherjee:** Writing – review & editing, Methodology, Formal analysis, Conceptualization. **Carola-Bibiane Schönlieb:** Supervision, Resources, Project administration, Conceptualization. **Alberto Redaelli:** Writing – review & editing, Supervision, Resources.

#### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

#### Acknowledgements

This research has received no external funding.






Fig. 16. WSS magnitude colormaps on the 3D aortic surface calculated from 4D flow data (row 1) and from SIREN results (row 2). WSS vectors on the 3D aortic surface calculated from 4D flow data (row 3) and from SIREN results (row 4). WSS magnitude colormaps on the unwrapped surface obtained by following [2] calculated from 4D flow data (row 5) and from SIREN results (row 6).




Fig. 17. Training and testing loss functions during training of a SIREN with 20 layers and 500 neurons per layer on data with *extreme* noise levels.

### References

- [1] D. Alblas, C. Brune, K.K. Yeung, J.M. Wolterink, Going off-grid: continuous implicit neural representations for 3D vascular modeling, in: *Statistical Atlases and Computational Models of the Heart. Regular and CMRxMotion Challenge Papers: 13th International Workshop, STACOM 2022, Held in Conjunction with MICCAI 2022, Singapore, September 18, 2022, Revised Selected Papers*, Springer, 2023, pp. 79–90.
- [2] L. Antiga, D.A. Steinman, Robust and objective decomposition and mapping of bifurcating vessels, *IEEE Trans. Med. Imaging* 23 (6) (2004) 704–713.
- [3] A. Bakhshinejad, A. Baghaie, A. Vali, D. Saloner, V.L. Rayz, R.M. D'Souza, Merging computational fluid dynamics and 4D flow MRI using proper orthogonal decomposition and ridge regression, *J. Biomech.* 58 (2017) 162–173.
- [4] A.J. Barker, M. Markl, J. Bürk, R. Lorenz, J. Bock, S. Bauer, J. Schulz-Menger, F. von Knobelsdorff-Brenkenhoff, Bicuspid aortic valve is associated with altered wall shear stress in the ascending aorta, *Circ. Cardiovasc. Imaging* 5 (4) (2012) 457–466.
- [5] R. Basri, M. Galun, A. Geifman, D. Jacobs, Y. Kasten, S. Kritchman, Frequency bias in neural networks for input of non-uniform density, in: *International Conference on Machine Learning*, PMLR, 2020, pp. 685–694.
- [6] M.M. Bissell, A.T. Hess, L. Biasioli, S.J. Glaze, M. Loudon, A. Pitcher, A. Davis, B. Prendergast, M. Markl, A.J. Barker, et al., Aortic dilation in bicuspid aortic valve disease: flow pattern is a major contributor and differs with valve fusion type, *Circ. Cardiovasc. Imaging* 6 (4) (2013) 499–507.
- [7] J. Bock, A. Frydrychowicz, R. Lorenz, D. Hirtler, A.J. Barker, K.M. Johnson, R. Arnold, H. Burkhardt, J. Hennig, M. Markl, In vivo noninvasive 4D pressure difference mapping in the human aorta: phantom comparison and application in healthy volunteers and patients, *Magn. Reson. Med.* 66 (4) (2011) 1079–1088.

- [8] J. Busch, D. Giese, L. Wissmann, S. Kozerke, Reconstruction of divergence-free velocity fields from cine 3D phase-contrast flow measurements, *Magn. Reson. Med.* 69 (1) (2013) 200–210.
- [9] M.F. Fathi, I. Perez-Raya, A. Baghaie, P. Berg, G. Janiga, A. Arzani, R.M. D'Souza, Super-resolution and denoising of 4D-flow MRI using physics-informed deep neural nets, *Comput. Methods Programs Biomed.* 197 (2020) 105,729.
- [10] E. Ferdian, A. Suinesiaputra, D.J. Dubowitz, D. Zhao, A. Wang, B. Cowan, A.A. Young, 4DFlowNet: super-resolution 4D flow MRI using deep learning and computational fluid dynamics, *Front. Phys.* (2020) 138.
- [11] S.W. Funke, M. Nordaas, Ø. Evju, M.S. Alnæs, K.A. Mardal, Variational data assimilation for transient blood flow simulations: cerebral aneurysms as an illustrative example, *Int. J. Numer. Methods Biomed. Eng.* 35 (1) (2019) e3152.
- [12] F. Gaidzik, D. Stucht, C. Roloff, O. Speck, D. Thévenin, G. Janiga, Transient flow prediction in an idealized aneurysm geometry using data assimilation, *Comput. Biol. Med.* 115 (2019) 103,507.
- [13] H. Ha, G.B. Kim, J. Kweon, S.J. Lee, Y.H. Kim, D.H. Lee, D.H. Yang, N. Kim, Hemodynamic measurement using four-dimensional phase-contrast MRI: quantification of hemodynamic parameters and clinical applications, *Korean J. Radiol.* 17 (4) (2016) 445–462.
- [14] M. Habibi, R.M. D'Souza, S.T. Dawson, A. Arzani, Integrating multi-fidelity blood flow data with reduced-order data assimilation, *Comput. Biol. Med.* 135 (2021) 104,566.
- [15] R. Izzo, D. Steinman, S. Manini, L. Antiga, The vascular modeling toolkit: a Python library for the analysis of tubular structures in medical images, *J. Open Sour. Softw.* 3 (25) (2018) 745.
- [16] J. Jiang, P. Kokeny, W. Ying, C. Magnano, R. Zivadinov, E.M. Haacke, Quantifying errors in flow measurement using phase contrast magnetic resonance imaging: comparison of several boundary detection methods, *Magn. Reson. Imaging* 33 (2) (2015) 185–193.
- [17] A. Kontogiannis, M.P. Juniper, Physics-informed compressed sensing for PC-MRI: an inverse Navier-Stokes problem, *IEEE Trans. Image Process.* 32 (2022) 281–294.
- [18] M. Markl, A. Frydrychowicz, S. Kozerke, M. Hope, O. Wieben, 4D flow MRI, *J. Magn. Reson. Imaging* 36 (5) (2012) 1015–1036.
- [19] J.N. Martel, D.B. Lindell, C.Z. Lin, E.R. Chan, M. Monteiro, G. Wetzstein, Acorn: adaptive coordinate networks for neural scene representation, *arXiv preprint, arXiv:2105.02788*, 2021.
- [20] F. Mut, R. Löhner, A. Chien, S. Tateshima, F. Viñuela, C. Putman, J.R. Cebra, Computational hemodynamics framework for the analysis of cerebral aneurysms, *Int. J. Numer. Methods Biomed. Eng.* 27 (6) (2011) 822–839.
- [21] F. Ong, M. Uecker, U. Tariq, A. Hsiao, M.T. Alley, S.S. Vasanawala, M. Lustig, Robust 4D flow denoising using divergence-free wavelet transform, *Magn. Reson. Med.* 73 (2) (2015) 828–842.
- [22] S. Petersson, P. Dyverfeldt, T. Ebbers, Assessment of the accuracy of MRI wall shear stress estimation using numerical simulations, *J. Magn. Reson. Imaging* 36 (1) (2012) 128–138.
- [23] F. Piatti, F. Sturla, M.M. Bissell, S. Pirola, M. Lombardi, I. Nesteruk, A. Della Corte, A.C. Redaelli, E. Votta, 4D flow analysis of BAV-related fluid-dynamic alterations: evidences of wall shear stress alterations in absence of clinically-relevant aortic anatomical remodeling, *Front. Physiol.* 8 (2017) 441.
- [24] N. Rahaman, A. Baratin, D. Arpit, F. Draxler, M. Lin, F. Hamprecht, Y. Bengio, A. Courville, On the spectral bias of neural networks, in: *International Conference on Machine Learning*, PMLR, 2019, pp. 5301–5310.
- [25] M. Raissi, P. Perdikaris, G.E. Karniadakis, Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, *J. Comput. Phys.* 378 (2019) 686–707.
- [26] G. Reiter, U. Reiter, G. Kovacs, H. Olschewski, M. Fuchsjaeger, Blood flow vortices along the main pulmonary artery measured with MR imaging for diagnosis of pulmonary hypertension, *Radiology* 275 (1) (2015) 71–79.
- [27] V.C. Rispoli, J.F. Nielsen, K.S. Nayak, J.L. Carvalho, Computational fluid dynamics simulations of blood flow regularized by 3D phase contrast MRI, *Biomed. Eng. Online* 14 (1) (2015) 1–23.
- [28] D.R. Rutkowski, A. Roldán-Alzate, K.M. Johnson, Enhancement of cerebrovascular 4D flow MRI velocity fields using machine learning and computational fluid dynamics simulation data, *Sci. Rep.* 11 (1) (2021) 1–11.
- [29] S. Saitta, B. Guo, S. Pirola, C. Menichini, D. Guo, Y. Shan, Z. Dong, X.Y. Xu, W. Fu, Qualitative and quantitative assessments of blood flow on tears in type B aortic dissection with different morphologies, *Front. Bioeng. Biotechnol.* (2021) 920.
- [30] S. Saitta, L. Maga, C. Armour, E. Votta, D.P. O'Regan, M.Y. Salmasi, T. Athanasiou, J.W. Weinsaft, X.Y. Xu, S. Pirola, A. Redaelli, Data-driven generation of 4D velocity profiles in the aneurysmal ascending aorta, *Comput. Methods Programs Biomed.* 233 (2023) 107,468, <https://doi.org/10.1016/j.cmpb.2023.107468>, <https://www.sciencedirect.com/science/article/pii/S0169260723001347>.
- [31] S. Saitta, S. Pirola, F. Piatti, E. Votta, F. Lucherini, F. Pluchinotta, M. Carminati, M. Lombardi, C. Geppert, F. Cuomo, et al., Evaluation of 4D flow MRI-based non-invasive pressure assessment in aortic coarctations, *J. Biomech.* 94 (2019) 13–21.
- [32] M.Y. Salmasi, S. Pirola, S. Sasidharan, S.M. Fischella, A. Redaelli, O.A. Jarral, D.P. O'Regan, A.Y. Oo, J.E. Moore Jr., X.Y. Xu, et al., High wall shear stress can predict wall degradation in ascending aortic aneurysms: an integrated biomechanics study, *Front. Bioeng. Biotechnol.* (2021) 935.
- [33] S. Shit, J. Zimmermann, I. Ezhov, J.C. Paetzold, A.F. Sanches, C. Pirkel, B. Menze, SR-flow: deep learning based super-resolution of 4D-flow MRI data, *Front. Artif. Intell.* 171 (2022).
- [34] V. Sitzmann, J. Martel, A. Bergman, D. Lindell, G. Wetzstein, Implicit neural representations with periodic activation functions, *Adv. Neural Inf. Process. Syst.* 33 (2020) 7462–7473.
- [35] J. Sotelo, L. Dux-Santoy, A. Guala, J. Rodríguez-Palomares, A. Evangelista, C. Sing-Long, J. Urbina, J. Mura, D.E. Hurtado, S. Uribe, 3D axial and circumferential wall shear stress from 4D flow MRI data using a finite element method and a laplacian approach, *Magn. Reson. Med.* 79 (5) (2018) 2816–2823.
- [36] M. Tancik, B. Mildenhall, T. Wang, D. Schmidt, P.P. Srinivasan, J.T. Barron, R. Ng, Learned initializations for optimizing coordinate-based neural representations, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 2846–2855.
- [37] M. Tancik, P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. Barron, R. Ng, Fourier features let networks learn high frequency functions in low dimensional domains, *Adv. Neural Inf. Process. Syst.* 33 (2020) 7537–7547.
- [38] J. Thewlis, D. Stevens, H. Power, D. Giddings, P. Gowland, M. Vloeberghs, 4-dimensional local radial basis function interpolation of large, uniformly spaced datasets, *Comput. Methods Programs Biomed.* 228 (2023) 107,235.
- [39] J. Töger, M.J. Zahr, N. Aristokleous, K. Markenroth Bloch, M. Carlsson, P.O. Persson, Blood flow imaging by optimal matching of computational fluid dynamics to 4D-flow data, *Magn. Reson. Med.* 84 (4) (2020) 2231–2245.
- [40] C. Trenti, M. Ziegler, N. Bjarnegård, T. Ebbers, M. Lindenberg, P. Dyverfeldt, Wall shear stress and relative residence time as potential risk factors for abdominal aortic aneurysms in males: a 4D flow cardiovascular magnetic resonance case–control study, *J. Cardiovasc. Magn. Reson.* 24 (1) (2022) 1–12.
- [41] D. Ulyanov, A. Vedaldi, V. Lempitsky, Deep image prior, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 9446–9454.
- [42] V. Vishnevskiy, J. Walheim, S. Kozerke, Deep variational network for rapid 4D flow MRI reconstruction, *Nat. Mach. Intell.* 2 (4) (2020) 228–235.
- [43] S. Voß, O. Beuing, G. Janiga, P. Berg, Multiple aneurysms anatomy challenge 2018 (MATCH)—phase Ib: effect of morphology on hemodynamics, *PLoS ONE* 14 (5) (2019) e0216,813.
- [44] J.M. Wolterink, J.C. Zwienenberg, C. Brune, Implicit neural representations for deformable image registration, in: *International Conference on Medical Imaging with Deep Learning*, PMLR, 2022, pp. 1349–1359.
- [45] P.A. Yushkevich, Y. Gao, G. Gerig, ITK-SNAP: an interactive tool for semi-automatic segmentation of multi-modality biomedical images, in: *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBC, IEEE*, 2016, pp. 3342–3345.