



Deposited via The University of Leeds.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/197135/>

Version: Accepted Version

---

**Proceedings Paper:**

Shone, F, Ravikumar, N, Lassila, T et al. (2023) Deep Physics-Informed Super-Resolution of Cardiac 4D-Flow MRI. In: Information Processing in Medical Imaging. Information Processing in Medical Imaging (IPMI 2023), 18-23 Jun 2023, San Carlos de Bariloche, Argentina. Lecture Notes in Computer Science. Springer, pp. 511-522. ISBN: 978-3-031-34047-5. ISSN: 0302-9743. EISSN: 1611-3349.

[https://doi.org/10.1007/978-3-031-34048-2\\_39](https://doi.org/10.1007/978-3-031-34048-2_39)

---

© 2023 The Author(s), under exclusive license to Springer Nature Switzerland AG. This version of the conference paper has been accepted for publication, after peer review (when applicable) and is subject to Springer Nature's AM terms of use (<https://www.springernature.com/gp/open-research/policies/accepted-manuscript-terms>), but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: [https://doi.org/10.1007/978-3-031-34048-2\\_39](https://doi.org/10.1007/978-3-031-34048-2_39)

**Reuse**

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

**Takedown**

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.

# Deep Physics-Informed Super-Resolution of Cardiac 4D-Flow MRI

Fergus Shone<sup>1,2,3,4</sup>, Nishant Ravikumar<sup>1,2</sup>, Toni Lassila<sup>1,2,3</sup>, Michael MacRaid<sup>1,2,3</sup>, Yongxing Wang<sup>1,2</sup>, Zeike A Taylor<sup>1,2</sup>, Peter Jimack<sup>3</sup>, Erica Dall’Armellina<sup>1,2,4\*</sup>, and Alejandro F Frangi<sup>1,2,3,4\*</sup>

<sup>1</sup> Centre for Computational Imaging and Simulation Technologies in Biomedicine, Schools of Computing, Mechanical Engineering and Medicine, University of Leeds, UK

`mm16f2s@leeds.ac.uk`

<sup>2</sup> National Institute for Health and Care Research (NIHR) Leeds Biomedical Research Centre (BRC), UK

<sup>3</sup> EPSRC Centre for Doctoral Training in Fluid Dynamics, School of Computing, University of Leeds, UK

<sup>4</sup> Leeds Institute of Cardiovascular and Metabolic Medicine, School of Medicine, University of Leeds, UK

**Abstract.** 4D-flow magnetic resonance imaging (MRI) provides non-invasive blood flow reconstructions in the heart. However, low spatio-temporal resolution and significant noise artefacts hamper the accuracy of derived haemodynamic quantities. We propose a physics-informed super-resolution approach to address these shortcomings and uncover hidden solution fields. We demonstrate the feasibility of the model through two synthetic studies generated using computational fluid dynamics. The Navier-Stokes equations and no-slip boundary condition on the endocardium are weakly enforced, regularising model predictions to accommodate network training without high-resolution labels. We show robustness to each type of data degradation, achieving normalised velocity RMSE values of under 16% at extreme spatial and temporal upsampling rates of  $16\times$  and  $10\times$  respectively, using a signal-to-noise ratio of 7.

**Keywords:** Physics-informed machine learning · 4D-flow MRI · Super-resolution.

## 1 Introduction

**4D-Flow MRI.** Phase-contrast magnetic resonance imaging (PC-MRI) is the modality of choice for quantifying cardiovascular blood flow across a range of target domains. For visualising cardiac flow, 4D-flow MRI [8] is the preferred option, due to its ability to provide time-resolved, 3D reconstructions of the velocity field. However, its direct application in clinic has been hampered by several shortcomings: low spatio-temporal resolution, significant noise artefacts

---

\* Joint senior authors.

and long scan times. Moreover, clinically relevant derived quantities like pressure, vorticity and wall shear stresses are not directly measured, and thus susceptible to effects of corruption in the velocity data. Thus, efforts have been made to increase the spatio-temporal resolution and denoise 4D-flow MRI data to improve the accuracy of predictive haemodynamic quantities [11,4,3].


**Super-Resolution of 4D-Flow MRI.** Unlike in structural MRI, super-resolution of 4D-flow MRI is an emerging field, with few publications to date. Machine learning approaches have dominated recent efforts, all of which have focused entirely on flow in the vasculature. Residual networks are used in [4] and [11], increasing the spatial resolution by factors of  $2\times$  and  $4\times$  respectively whilst denoising the data. Both approaches here used architectures designed for single image super-resolution, and therefore address only spatial upsampling. This not only precludes temporal super-resolution, but also ignores rich information in flow data at neighbouring time points, which could also improve spatial upsampling capabilities. Furthermore, both approaches require paired low- and high-resolution data during training, which is not readily available in 4D-flow MRI studies. In contrast, super-resolution in both space and time was achieved in [3] using a physics-informed neural network (PINN), without high-resolution target data. Upsampling factors of  $100\times$  and  $5\times$  in space and time respectively were reported. However, the network architecture and hyperparameters used in [3] are not capable of capturing complex, small-scale flow features, like those present in cardiac flow. The model used in [3] belongs to a class of supervised machine learning methods called PINNs [10] which are designed to model physical systems for which only sparse and potentially noisy data are available. Weak imposition of symmetries, physical laws and other *a priori* domain knowledge via a multi-component loss function provides regularisation during network training to heavily restrict the space of possible solutions, and render this small-data setting feasible.

**Our Contributions.** We present the first application of PINNs to fluid flow problems in moving domains, and the first instance of flow super-resolution in the left ventricle (LV). We demonstrate the feasibility of our model using two synthetic cases, generated using computational fluid dynamics (CFD): a 2D idealised ventricle, with a focus on exploring network parameter configurations, including a comparative study with various activation functions; and a 3D patient-specific LV geometry, for which we complete a further network hyperparameter study. In both cases, we evaluate robustness to different levels of image degradation, testing performance at different spatial/temporal downsampling rates and noise levels.

## 2 Methods

**PINN Model.** The proposed PINN model uses a fully connected neural network to approximate a function mapping spatio-temporal coordinates on a do-

main to corresponding velocity and pressure fields, as seen in Fig. 1. Predicted solutions are constrained through a multi-component loss function involving low-resolution measurements, boundary conditions, and known physical laws. Namely: 4D-flow MRI data, the no-slip condition on the endocardium, and residuals of the Navier-Stokes equations, respectively. Unlike purely data-driven machine learning approaches, the physics-based constraints used in PINNs restrict the space of possible solutions by penalising non-physical predictions, facilitating efficient training with sparse and noisy data whilst ensuring predicted solution fields obey the underlying physics. With this model, we aim to only use information that would be present in a real cardiac 4D-flow MRI study, where the boundary motion of the endocardium is extracted through segmentation and registration of structural cine-MR images.



**Fig. 1.** Basic overview of the physics-informed neural network used in this paper, displaying network inputs and outputs, and the three loss components.

In large vessels and cardiac chambers, blood is considered a Newtonian, incompressible fluid. Therefore, the flow is governed by the incompressible Navier-Stokes equations. After nondimensionalising our input and output variables using characteristic length and velocity scales [7] and standardising our input variables, they are expressed as

$$\mathbf{F}_{1-3} = \frac{1}{\sigma_{t^*}} \partial_t \hat{\mathbf{u}} + (\hat{\mathbf{u}} \cdot \hat{\nabla}) \hat{\mathbf{u}} + \hat{\nabla} \hat{p} - \frac{1}{\text{Re}} \hat{\nabla}^2 \hat{\mathbf{u}} = 0 \quad (1)$$

$$F_4 = \hat{\nabla} \cdot \hat{\mathbf{u}} = 0, \quad (2)$$

where

$$\hat{\nabla} = \left( \frac{1}{\sigma_{x_1^*}} \partial_{\hat{x}_1}, \frac{1}{\sigma_{x_2^*}} \partial_{\hat{x}_2}, \frac{1}{\sigma_{x_3^*}} \partial_{\hat{x}_3} \right)^T, \quad \hat{\nabla}^2 = \frac{1}{\sigma_{x_1^*}^2} \partial_{\hat{x}_1}^2 + \frac{1}{\sigma_{x_2^*}^2} \partial_{\hat{x}_2}^2 + \frac{1}{\sigma_{x_3^*}^2} \partial_{\hat{x}_3}^2, \quad (3)$$

for Reynolds number  $\text{Re} = \rho UL/\mu$ , with dimensionless velocity  $\hat{\mathbf{u}} = \mathbf{u}/U$  and pressure  $\hat{p} = p/\rho U^2$ , density  $\rho$  and dynamic viscosity  $\mu$ .  $\sigma_i$  values are standard deviations of each input variable. We assume  $\rho = 1066 \text{ kg m}^{-3}$  and  $\mu = 0.0035 \text{ Pa s}$  [7]. For wall velocity  $\mathbf{u}^w$ , the no-slip boundary condition is enforced as  $U \hat{\mathbf{u}} = \mathbf{u}^w$ .

**Loss Function.** The loss function consists of three components: *Data loss*,  $\mathcal{L}_{data}$ , in which the error between velocity predictions and 4D-flow MRI measurements is minimised, *PDE (partial differential equation) loss*,  $\mathcal{L}_{PDE}$ , containing residuals of the Navier-Stokes equations, and *BC (boundary condition) loss*,  $\mathcal{L}_{BC}$ , where predictions on the domain wall are constrained to obey the no-slip condition. In all cases, the  $L_2$  loss is used, as per [10]. The total loss to be minimised during training is given by

$$\mathcal{L} = \mathcal{L}_{PDE} + \alpha\mathcal{L}_{data} + \beta\mathcal{L}_{BC}, \quad (4)$$

where

$$\mathcal{L}_{PDE} = \frac{1}{m} \sum_{k=1}^m \sum_{i=1}^4 \|F_i(\hat{\mathbf{x}}_k, \hat{t}_k, \hat{\mathbf{u}}_k, \hat{p}_k)\|^2 \quad (5)$$

$$\mathcal{L}_{data} = \frac{1}{n} \sum_{k=1}^n \|(U\hat{\mathbf{u}}(\hat{\mathbf{x}}_k, \hat{t}_k) - \mathbf{u}_k)/U\|^2 \quad (6)$$

$$\mathcal{L}_{BC} = \frac{1}{p} \sum_{k=1}^p \|(U\hat{\mathbf{u}}(\hat{\mathbf{x}}_k, \hat{t}_k) - \mathbf{u}_k^w)/U\|^2. \quad (7)$$

Here,  $m$ ,  $n$  and  $p$  are numbers of collocation, data and wall training points, respectively,  $\mathbf{u}_k$  are velocity measurements, and  $\mathbf{u}_k^w$  are wall velocity measurements.

In equation 4,  $\alpha$  and  $\beta$  are weighting coefficients selected dynamically during training, as per [5]. Incorporating a loss weighting scheme such as this is highly important when using PINNs to balance the gradient contributions from each loss component during training. At training iteration  $k$ , we calculate

$$\hat{\alpha}^{k+1} = \frac{|\overline{\nabla_{\theta}\mathcal{L}_{PDE}}|}{|\overline{\nabla_{\theta}\mathcal{L}_{data}}|}, \quad \hat{\beta}^{k+1} = \frac{|\overline{\nabla_{\theta}\mathcal{L}_{PDE}}|}{|\overline{\nabla_{\theta}\mathcal{L}_{BC}}|}, \quad (8)$$

where  $\nabla_{\theta}\mathcal{L}_i$  is the gradient of loss component  $i$  w.r.t. the network weights,  $\theta$ . This operation is performed layer-wise and we take the mean value. We then calculate weights for epoch  $k+1$  as

$$\alpha^{k+1} = (1 - \lambda)\alpha^k + \lambda\hat{\alpha}^{k+1}, \quad \beta^{k+1} = (1 - \lambda)\beta^k + \lambda\hat{\beta}^{k+1}, \quad (9)$$

for constant  $\lambda$ , chosen to be 0.1 as per [5].

**Network Architecture.** To determine optimal network architecture and hyperparameter choices, we performed individual ablation studies for the 2D and 3D cases. Details of these studies can be found in the Experiments section. The networks used in both cases are fully connected, where hidden layers use Siren activation functions [12] and the output layer uses a linear activation. For both cases, the optimal number of layers was 9. For the 2D case, 900 neurons per layer with a dropout rate of 0.55 performed best, whereas in the 3D case these

values were 750 and 0.35, respectively. For hidden layer  $i$ , the Siren activation is formulated as

$$\phi_i(x_i) = \sin(\mathbf{W}_i x_i + \mathbf{b}_i), \quad (10)$$

for input  $x_i$ , weights  $\mathbf{W}_i$  and biases  $\mathbf{b}_i$ , where weights in each hidden layer,  $w_i$ , are drawn from  $w_i \sim \mathcal{U}(-\sqrt{6/n}, \sqrt{6/n})$  at initialisation. This activation function is designed specifically to capture high-frequency solution information and model higher-order derivatives, both of which are relevant to this study. Further details about this activation function can be found in [12].

ADAM [6] is used as the optimiser with an initial learning rate of  $1 \times 10^{-5}$  for the 2D case, and  $1 \times 10^{-4}$  for the 3D case, which is selected to be the highest possible value without divergence. Plateau-based annealing is used, where the learning rate decays by a factor of 0.1 if the validation loss plateaus, within a tolerance of  $10^{-6}$ . The annealing initiates after 5 epochs, and has a cool down period of 3 epochs following a decay action. All code is written in Python, primarily using the package TensorFlow 2 [1].

**Error Metrics.** We evaluate performance using max-normalised root mean square error (NRMSE) between predicted solution fields and full-resolution ground truth fields, with velocity NRMSE given by:

$$\text{vNRMSE} := \frac{1}{\max|\mathbf{u}|} \sqrt{\frac{1}{N} \sum_{k=1}^N (\mathbf{u}_{pred} - \mathbf{u})_k^2},$$

where  $\mathbf{u}_{pred} = U\hat{\mathbf{u}}$ , and  $\mathbf{u}$  is the ground truth velocity, at  $N$  data locations. We also calculate pressure NRMSE by zero-centring the mean value for prediction and ground truth data at each time step. The pressure NRMSE then assumes the same form as velocity NRMSE.

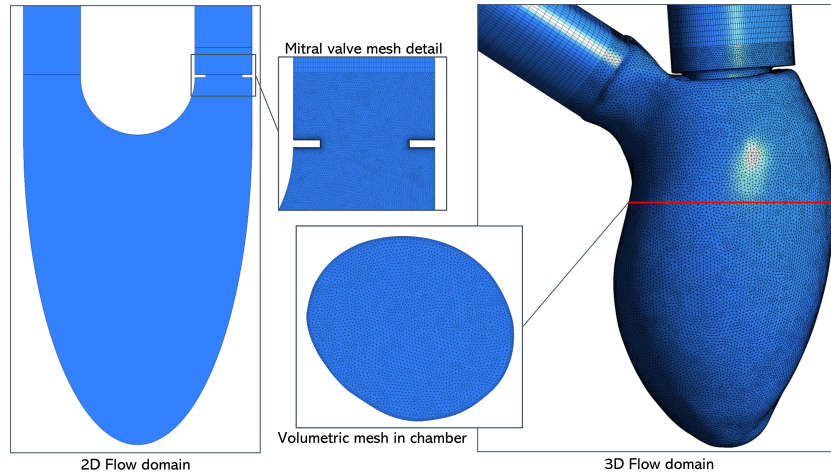
### 3 Experiments and Results

In this section, we describe our two experiments used to demonstrate the feasibility of our model. Both cases are synthetic, with solution ground truth data generated from CFD simulations. Use of synthetic data is of particular importance in this study, as such higher resolution ground truth data are not available for *in vivo* cases.

#### 3.1 Case 1: 2D Idealised Ventricle

**Computational Fluid Dynamics.** Case 1 consists of a 2D idealised ventricle, as seen on the left in Fig. 2. The simulation was computed in ANSYS Fluent (ANSYS Inc., Canonsburg, PA). Boundary motion across diastole and systole was prescribed analytically with a sinusoidal function, driving the flow in both

phases. The period of one complete cardiac cycle was 0.5 seconds, resulting in a peak  $Re = 5600$ , in line with expectations for true LV flow. In diastole, the inlet was opened and the outlet closed, and vice versa in systole. Two small protrusions were placed at the inlet to represent the open leaflets of the mitral valve. Zero-normal stress conditions were applied to the inlet and outlet when opened, with the no-slip condition applied on the walls and closed openings. We used a highly-resolved mesh of  $\sim 330k$  elements and error-based adaptive time-stepping to capture flow features across a range of scales, and to emulate complex flow patterns observed in a real LV. The minimum time step was  $10^{-6}$  seconds. Adaptive meshing was used with diffusion-based smoothing, remeshing every 3 time steps based on minimum and maximum length scales determined by the initial mesh sizing.

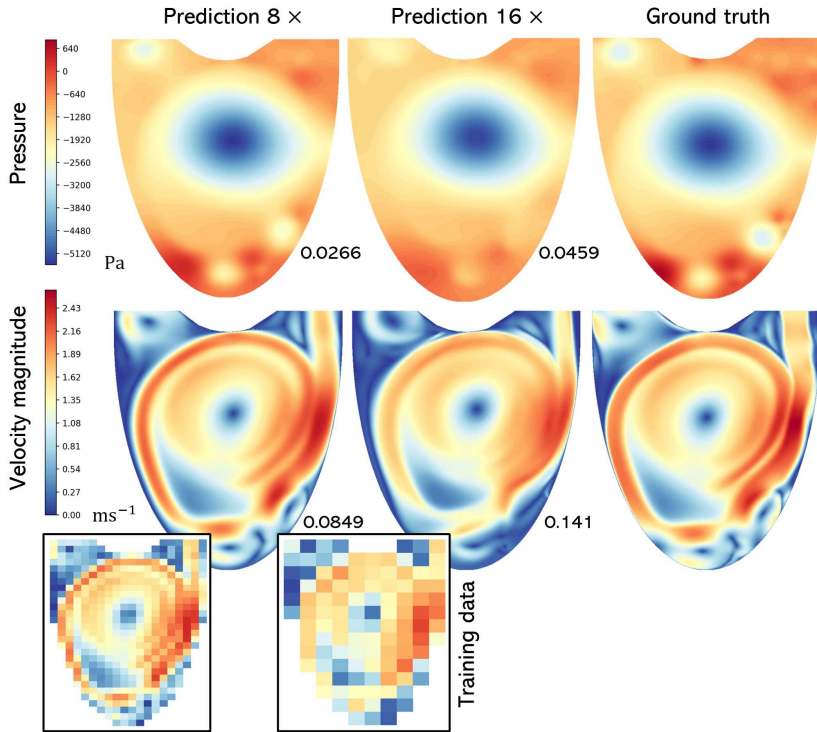


**Fig. 2.** Geometry and mesh details for the 2D idealised left ventricle (left) and 3D patient-specific left ventricle (right). In end-diastole, the 2D mesh consisted of  $\sim 0.33$  million elements, whilst the 3D mesh consisted of  $\sim 1.3$  million elements.

**Training Data Generation.** Temporal downsampling of the data was performed by simply removing requisite frames from the cardiac cycle. The spatial downsampling process was designed to reflect real 4D-flow MRI acquisition processes: 1) CFD solution data were interpolated on a high-resolution Cartesian grid at each time step; 2) a square region of interest (ROI) was extracted, removing inflow and outflow channels; 3) four-point balanced encoding [9] was applied to the velocity data, with the encoding velocity set to 10% higher than the maximum velocity; 4) velocity ‘images’ were converted to  $k$ -space images using the fast Fourier transform (FFT); and 5) spatial downsampling was achieved by cropping the outer edges of the  $k$ -space images, with the effect of truncating high-frequency modes [4]. Zero-mean Gaussian noise was added in the frequency

domain to both real and imaginary signals, matching the true noise distribution of 4D-Flow MRI data [4,3]. To control noise levels, we altered the standard deviation, which was based on a signal power calculated using the encoding velocity. Finally, inverse FFT (IFFT) and balanced four-point decoding were applied to recover the downsampled velocity data in the spatial domain. The maximum downsampling rates and noise levels simulated exceed what would be expected in real 4D-flow MRI data to demonstrate the effective range of the model.

In studies with real data, the boundary motion and collocation point clouds are generated using segmentation and registration of structural MRI images acquired alongside the 4D-flow MRI data. However, in the case of the synthetic studies discussed in this paper, we simply used the CFD results at the boundary to mask the flow domain and provide the wall coordinates and associated velocities. The collocation point cloud was chosen to be the cell centres of the underlying CFD mesh, since this was already highly resolved and evenly distributed in space and time. This produced a set of  $\sim 6,000,000$  collocation points throughout space and time.



**Fig. 3.** 2D idealised ventricle: pressure (top row) and velocity magnitude (second row) contour plots in late diastole, using noise-free training data, temporal upsampling rate of  $5\times$ , and spatial upsampling rates of  $8\times$  (left) and  $16\times$  (centre), alongside the ground truth (right). Inserts show the resolution of training data. The values reported next to each figure are NRMSE for the given case.

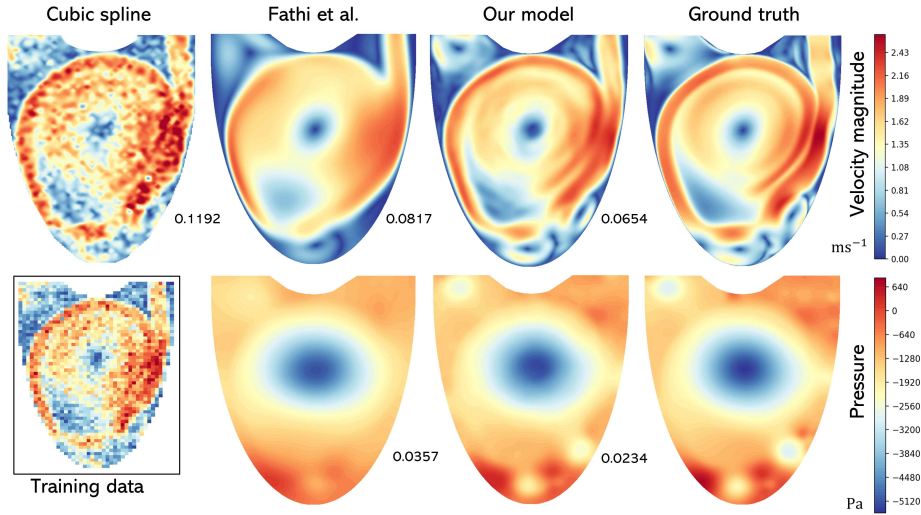
The complete synthetic 4D-flow MRI dataset was split 80:20 for training and validation. We constructed our test set using 10% of the ground truth CFD data evenly distributed in space and time. We confirmed that batch size has a minimal effect on accuracy, and thus selected the maximum possible batch size for each case for computational efficiency. Since the size of each training set differs significantly, we continuously looped through the batches of the smaller sets without shuffling them, while the largest set completed the epoch. Then, all training sets were shuffled before generating new batches.

**Results.** In Fig. 3, we see network predictions in late diastole at spatial upsampling rates of  $8\times$  and  $16\times$ , using a temporal upsampling rate of  $5\times$ , and zero noise. To ensure pressure colour bars are matched, we shifted the predicted pressure values by the difference in predicted and ground truth means. This is required because pressure predictions are only unique up to a constant, given that we have no pressure measurements and pressure appears only as a derivative in the Navier-Stokes equations. We see good qualitative agreement in both velocity and pressure fields, even at the higher spatial upsampling rate, with much of the fine-scale flow captured well. The results far outperform the cubic spline interpolation used for comparison, which achieves vNRMSE values of 0.115 and 0.181 for the two spatial upsampling rates. Pressure NRMSE remains low (0.0459) even at  $16\times$  upsampling, while the velocity NRMSE is rather higher, at 0.141. This could partly be due to small misalignment of high-velocity regions in the flow, e.g. in the left region of the ventricle.

We assess the denoising capabilities of our model in Fig. 4, where we add noise with a standard deviation of 0.15 to our 4D-flow MRI training data, and compare the performance against two competing methods: cubic spline interpolation and the PINN configuration used in Section 4.2 of [3]. Qualitatively, our model shows robustness to noise and is able to capture most of the solution features present in the ground truth data. Conversely, the cubic spline algorithm is heavily corrupted by the presence of noise and is clearly unsuitable in this context, while the PINN model used in [3] is able to remove the data noise but eradicates much of the finer flow detail, which is also shown quantitatively with increased NRMSE values. In the presence of noise, the spatial upsampling capabilities of our model are more limited, where NRMSE values at the spatial upsampling rate of  $4\times$  with noise are comparable to those at  $8\times$  with noise-free data. However, in this work we have not deployed any specific component to negate the effects of noise, such as label smoothing, and the inclusion of such components should improve robustness to these effects.

### 3.2 Case 2: Patient-Specific Left Ventricle

**CFD Setup.** As before, results for this case were produced using ANSYS Fluent. The mean of the principle component analysis (PCA) shape model in [13] was used as a 3D LV geometry. Inflow and outflow channels were added at mitral and aortic openings to ensure fully developed inlet and outlet flows, which were

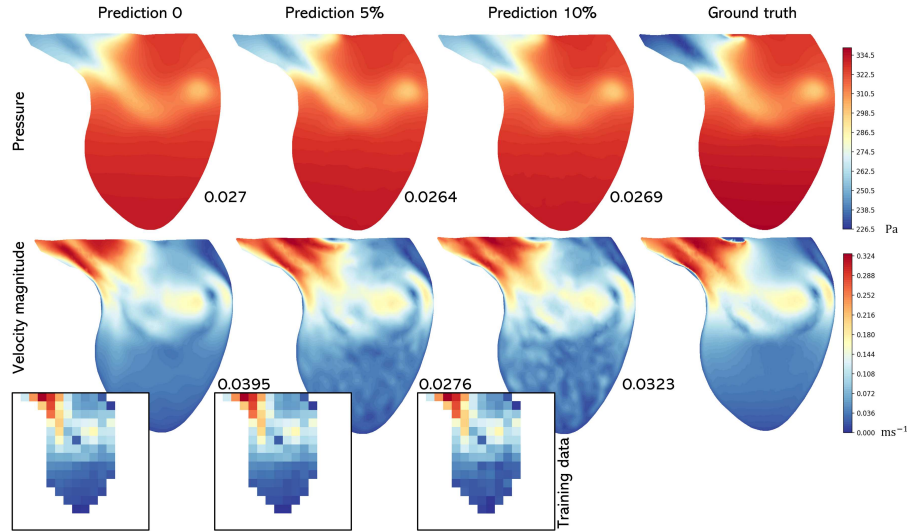


**Fig. 4.** 2D idealised ventricle: velocity magnitude (top row) and pressure (bottom row) contour plots in late diastole, comparing our model (middle right) with that used in [3] (middle left) and cubic spline (left). Here we have used noisy (standard deviation: 15%) training data with temporal and spatial upsampling rates of  $5\times$  and  $4\times$  respectively. Insert displays the resolution of training data. Again, reported values are NRMSE.

again modelled using zero-normal stress conditions. Aortic and mitral valves opened and closed at corresponding times in the cardiac cycle, whose period was set to 1 second. For simplicity, we assumed either fully open or closed valves, with no opening/closing phase in between. In diastole, we approximated the open mitral valve shape using the mitral valve plane generated in [13], with some intrusion of the valve leaflets into the flow domain. Endocardial motion was sinusoidal as for the 2D case, with the ventricular base fixed and lower regions toward the apex moving, and with the same time-stepping configuration. In end-diastole, the mesh consisted of  $\sim 1.3\text{M}$  elements. Fig. 2 (right) depicts the geometry and mesh.

**Training Data Generation.** The downsampling procedure was similar to the 2D case, with the corresponding 3D algorithms used for FFT and IFFT operations. The ROI was selected to remove most of the inflow and outflow channels, and the CFD mesh cell centres were again used as coordinates for the collocation point cloud, producing  $\sim 25\text{M}$  points in space and time.

**Results.** A selection of results is shown in Fig. 5, where we compare the effect of varying noise levels at temporal and spatial upsampling rates of  $5\times$  and  $8\times$  respectively. Pressure NRMSE values are mean-normalised, as the maximum pressures in the ground truth data were unrealistically large due to some of the outflow channel being included in the dataset. Visually, noise has a clear



**Fig. 5.** 3D patient-specific ventricle: pressure (top row) and velocity magnitude (second row) contour plots across central x-z plane in early systole, using training data at a temporal upsampling rate of  $5\times$  and spatial upsampling rate of  $8\times$ , in the presence of noise with a standard deviation of 5% (centre left), 10% (centre right) and no noise (left), alongside the ground truth (right). Inserts display resolution of training data. Again, reported values are NRMSE, although the pressure NRMSE values are normalised using the mean instead of the maximum.

effect in low velocity regions toward the apex, but flow structures toward the base and valve region are well captured. This occurs because velocity scales with signal intensity in 4D-flow MRI, and thus higher velocity voxels have a higher signal-to-noise ratio. In contrast, pressure predictions appear far more robust to noise levels, reflected qualitatively and quantitatively in the contour plots and NRMSE values. This again supports the hypothesis that the PDE regularisation component is robust to uncertainty generated from noise in the velocity measurements. Some discrepancy is apparent in the low pressure regions near the aortic valve. However, this could be due to inconsistency in the ground truth data. Part of this region belongs to the outflow channel on the other side of the aortic valve plane, which experiences unrealistically extreme pressures when the valve is closed.

### 3.3 Hyperparameter Optimisation Studies

To perform our hyperparameter optimisation studies we used the Bayesian tree-structured parzen estimator algorithm from the Python package Hyperopt [2]. The ranges set in the search space were based on maximum and minimum values associated with PINN models seen in the literature. Complete search spaces for both studies can be found in Tab. 1. We performed 20 complete training runs per hyperparameter option, resulting in 140 training runs in total. For the 3D

case, we performed a slightly restricted study based on the assumption that the Siren activation function would remain the optimal choice. Although optimal network depth remained the same for both 2D and 3D cases, the optimal width was found to be smaller for the 3D case. This, perhaps unexpectedly, suggests that network size scales more closely with complexity of flow rather than number of spatial dimensions, with results in the 2D case exhibiting more complicated flow patterns.

**Table 1.** Search spaces defined for the 2D and 3D hyperparameter optimisation studies

Case	Hidden layers	Neurons per layer	Dropout rate	Activation function
2D	5-12	100-1000	0-0.7	Tanh, Swish, Siren
3D	5-12	400-1000	0-0.7	-

## 4 Discussion and Conclusion

In this work, we propose a PINN-based method for super-resolution of 4D-flow MRI data in the LV and demonstrate its feasibility through two synthetic studies. The two cases presented were constructed using specific modelling and downsampling procedures to ensure that the data generated were representative of true cardiac 4D-flow MRI. We performed individual network ablation studies for both 2D and 3D cases, and then evaluated model performance at varying levels of data degradation, demonstrating robustness of the model to both noise artefacts and low spatio-temporal resolution at levels of data corruption beyond what would be expected in real cases. We found greater robustness to image degradation in the 3D case, although this is likely due to the simpler flow patterns simulated in the ground truth data, with a longer cardiac cycle and relatively coarser mesh used to produce the CFD results. In the presence of significant levels of noise, the spatial upsampling capabilities of the model in the 2D case are evidently restricted, however, the inclusion of specific denoising components could alleviate these issues. In both cases, pressure predictions were shown to be less affected by the presence of noise in the velocity data.

The capability of the model to super-resolve in both space and time, whilst also uncovering the hidden pressure fields without the use of any additional data, presents a clear advantage over previous 4D-flow MRI super-resolution approaches that are based on single image super-resolution models. Further, during training the model does not rely on the availability of paired low- and high-resolution data, which are not readily available in most 4D-flow MRI studies. The main disadvantage when compared with these methods is the inability of our model, and PINNs in general, to generalise across cases, meaning the network must be re-trained each time it is presented with a new patient or volunteer. However, accelerating training cycles using transfer learning or similar techniques could be a useful avenue of investigation here to compensate for the inherent lack of generalisation capabilities.

**Acknowledgements.** This work was partially supported by the EPSRC Centre for Doctoral Training in Fluid Dynamics (EP/L01615X/1) and the Royal Academy of Engineering Chair in Emerging Technologies (CiET1919/19). The computational work was undertaken on the UK National Tier-2 high performance computing service JADE-2 (EP/T022205/1).

## References

1. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado: TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems. 12th USENIX Symposium on OSDI (2016)
2. Bergstra, J., Yamins, D., Cox, D.D.: Making a Science of Model Search: Hyperparameter Optimization in Hundreds of Dimensions for Vision Architectures. *Comput. Sci. Discov.* **8** (2015)
3. 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. *Comput Methods Programs Biomed* **197** (2020)
4. Ferdian, E., Suinesiaputra, A., Dubowitz, D.J., Zhao, D., Wang, A., Cowan, B., Young, A.A.: 4DFlowNet: Super-Resolution 4D Flow MRI Using Deep Learning and Computational Fluid Dynamics. *Front. Phys.* **8**(183) (2020)
5. Jin, X., Cai, S., Li, H., Karniadakis, G.E.: NSFnets (Navier-Stokes flow nets): Physics-informed neural networks for the incompressible Navier-Stokes equations. *J. Comput. Phys.* **426** (2021)
6. Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization. 3rd ICLR (2015)
7. 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. *Comput Methods Appl Mech Eng* **358** (2020)
8. Markl, M., Frydrychowicz, A., Kozerke, S., Hope, M., Wieben, O.: 4D flow MRI. *J. Magn. Reson. Imaging* **36**(5) (2012)
9. Pelc, N.J., Bernstein, M.A., Shimakawa, A., Glover, G.H.: Encoding Strategies for Three-Direction Phase-Contrast MR Imaging of Flow'. *J. Magn. Reson. Imaging* (1991)
10. Raissi, M., Perdikaris, P., Karniadakis, G.E.: Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving non-linear partial differential equations. *Journal of Computational Physics* **378** (2019)
11. Shit, S., Zimmermann, J.: SRflow: Deep learning based super-resolution of 4D-flow MRI data. *Front. Artif. Intell.* (2022)
12. Sitzmann, V., Martel, J.N.P., Bergman, A.W., Lindell, D.B., Wetzstein, G.: Implicit Neural Representations with Periodic Activation Functions. 34th Conference on NeurIPS (2020)
13. Xia, Y., Chen, X., Ravikumar, N., Kelly, C., Attar, R., Aung, N., Neubauer, S., Petersen, S.E., Frangi, A.F.: Automatic 3D+t four-chamber CMR quantification of the UK biobank: integrating imaging and non-imaging data priors at scale. *Med. Image Anal.* **80** (2022)