ORIGINAL PAPER

Super-resolution 4D flow MRI to quantify aortic regurgitation using computational fluid dynamics and deep learning

Derek Long[1] · Cameron McMurdo[1] · Edward Ferdian[2] · Charlène A. Mauger[3] · David Marlevi[4,5] · Martyn P. Nash[1,3] · Alistair A. Young[2,6]

Abstract

Changes in cardiovascular hemodynamics are closely related to the development of aortic regurgitation (AR), a type of valvular heart disease. Metrics derived from blood flows are used to indicate AR onset and evaluate its severity. These metrics can be non-invasively obtained using four-dimensional (4D) flow magnetic resonance imaging (MRI), where accuracy is primarily dependent on spatial resolution. However, insufficient resolution often results from limitations in 4D flow MRI and complex aortic regurgitation hemodynamics. To address this, computational fluid dynamics simulations were transformed into synthetic 4D flow MRI data and used to train a variety of neural networks. These networks generated super-resolution, full-field phase images with an upsample factor of 4. Results showed decreased velocity error, high structural similarity scores, and improved learning capabilities from previous work. Further validation was performed on two sets of in vivo 4D flow MRI data and demonstrated success in de-noising flow images. This approach presents an opportunity to comprehensively analyse AR hemodynamics in a non-invasive manner.

Keywords Super-resolution · Four-dimensional flow magnetic resonance imaging (4D flow MRI) · Aortic regurgitation · Computational fluid dynamics (CFD) · Deep learning

Introduction

Aortic valve regurgitation, or aortic regurgitation (AR), is a common type of valvular heart disease where the aortic valve does not close properly, causing reflux of blood from the aorta into the left ventricle [1]. This reflux of blood is known as the regurgitant jet. Diagnosis and severity of AR is determined by evaluation of flow metrics, for example peak velocity, pressure drop, and regurgitant volume.

Cardiovascular four-dimensional (4D) flow magnetic resonance imaging (MRI) is a novel imaging technique to quantify full-field blood flow velocities, providing a three-dimensional (3D) velocity field across a region of interest throughout the cardiac cycle. Currently, due to the small

✉ Derek Long
  dlon450@aucklanduni.ac.nz

Cameron McMurdo
  cmcm658@aucklanduni.ac.nz

Edward Ferdian
  e.ferdian@auckland.ac.nz

Charlène A. Mauger
  c.mauger@auckland.ac.nz

David Marlevi
  marlevi@mit.edu

Martyn P. Nash
  martyn.nash@auckland.ac.nz

Alistair A. Young
  a.young@auckland.ac.nz

1   Department of Engineering Science, University of Auckland, Auckland, New Zealand

2   Department of Anatomy and Medical Imaging, University of Auckland, Auckland, New Zealand

3   Auckland Bioengineering Institute, University of Auckland, Auckland, New Zealand

4   Institute for Medical Engineering and Science, Massachusetts Institute of Technology, Cambridge, MA, USA

5   Department of Molecular Medicine and Surgery, Karolinska Institutet, Solna, Sweden

6   Department of Biomedical Engineering, King's College London, London, UK

width of the regurgitant jet and the limited spatiotemporal resolution of 4D flow MRI, it fails to accurately capture the complex hemodynamics of AR. The combination of computational fluid dynamics (CFD) and deep learning with 4D flow MRI will help to generate higher resolution images and recover hemodynamic parameters lost in current MRI images, with a target upsample factor of 4. This work will extend what has already been completed in 4DFlowNet [2] by using a wider range of flow characteristics to mimic AR, improving the data augmentation steps, and enhancing the artificial neural network with newer architecture structures.

AR occurs when the aortic valve does not close properly, causing blood to flow back into the left ventricle from the aorta. This forces the heart to work harder and pump more blood to the aorta, which can cause further heart problems in the future. AR also has varying levels of intensity, from trace or mild through moderate to severe [1]. Acute AR is considered a medical emergency as it can cause severe pulmonary edema and hypotension, that is, excess fluid in the lungs and low blood pressure, respectively. Patients with AR are monitored yearly with echocardiography to decide whether replacement of the aortic valve is necessary [3]. Flow metrics such as peak velocity and pressure drop are typically calculated non-invasively using 2D velocities from 2D Doppler echocardiography [1]. Due to limited information available in 2D, this method is known to overestimate pressure drops [4].

4D flow MRI is an established imaging technique that captures the temporal changes of 3D blood flow patterns within individual vascular structures [5, 6] and has proven promising for quantifying AR in clinical practice [7–10]. Velocities of blood particles are encoded in the phase of the MRI signal while the anatomy is visualised from the signal's magnitude [11]. However, 4D flow has several limitations, such as low spatiotemporal resolution, long scan time, and low signal-to-noise (SNR) ratio [12], which makes its clinical application to AR difficult. With a spatial resolution between 1.0 and 3.5 mm [6], details on the narrowest part of the jet cannot be captured as its width is typically much smaller than 3 mm [13] for mild AR. Therefore, spatial resolution is the biggest limitation in 4D flow MRI.

Deep learning [14] has had a significant impact in many scientific sectors, and is highly relevant in the field of medical imaging [15]. Advances in super-resolution image reconstruction [16] to obtain high-resolution (HR) images from low-resolution (LR) observations are increasingly being adopted for MRI with a deep learning-based approach [17, 18]. This approach is preferred as it not only has an advantage in spatial resolution quality over conventional super-resolution techniques [19], but also successfully denoises flow images [20].

The combination of 4D flow MRI with deep learning has been explored in multiple ways to increase resolution and provide more accurate estimates of physical quantities [2, 21, 22]. However, there are several limitations involved in this approach. Primarily, these have been related to insufficient data [2, 22], which is due to the requirement of paired LR and HR MR images. This can be difficult to obtain as HR MRI takes long scanning times and is subject to motion artifacts [2, 22, 23]. As an alternative, CFD models have been used to simulate 4D flow MRI as ground truth HR images, which are then downsampled to LR images [2, 22]. Other limitations include unstable and non-robust network architectures [21], which describe the organisational structure of the network's layers. The architectures play a significant role in the performance of the deep learning algorithm, as well as ignoring phase/velocity aliasing error [2, 20]. The aliasing error here refers to aliasing from having a velocity encoding, VENC [24] that is too low [25] rather than other types of MRI spatial aliasing which have been explored previously and reduced [26–28]. Note that VENC is an MR parameter to adjust the maximum velocity corresponding to a 360° phase shift in the data.

Recent development in object detection has proven significant in advancing ANN architecture [25], and appears to be widely used in many medical imaging applications [15]. Examples include residual blocks [29], dense blocks [30], and cross stage partial blocks [31], which show promise to increase network capacity mitigating degradation and memory utilisation issues.

Methods

Data generation

Modelling of the aortic valve was done using ANSYS 2021 R1, with CFX chosen for CFD simulations, and was an iterative process to determine the best parameters and options that would increase efficiency and accuracy.

Geometries

The design of the problem geometry focused on simplicity, over replicating a real-life aortic valve during regurgitation, to capture the flow features of the regurgitant jet, which are not captured in the 4D flow data at typical resolution, and not handled well by 4DFlowNet. The jet characteristics are modelled using a variety of eccentricities and angulations to capture the variation in jet properties, giving the additional flow information required for training of the flow super-resolution network. The basic geometry is radially symmetrical and shaped like a 3D cylinder with a constricted section part way along the length, resembling a Venturi tube. The constricted section represents the gap in the aortic valve

Fig. 1 Sketches of the basic (top), angled (middle), and offset (bottom) constricted sections, respectively
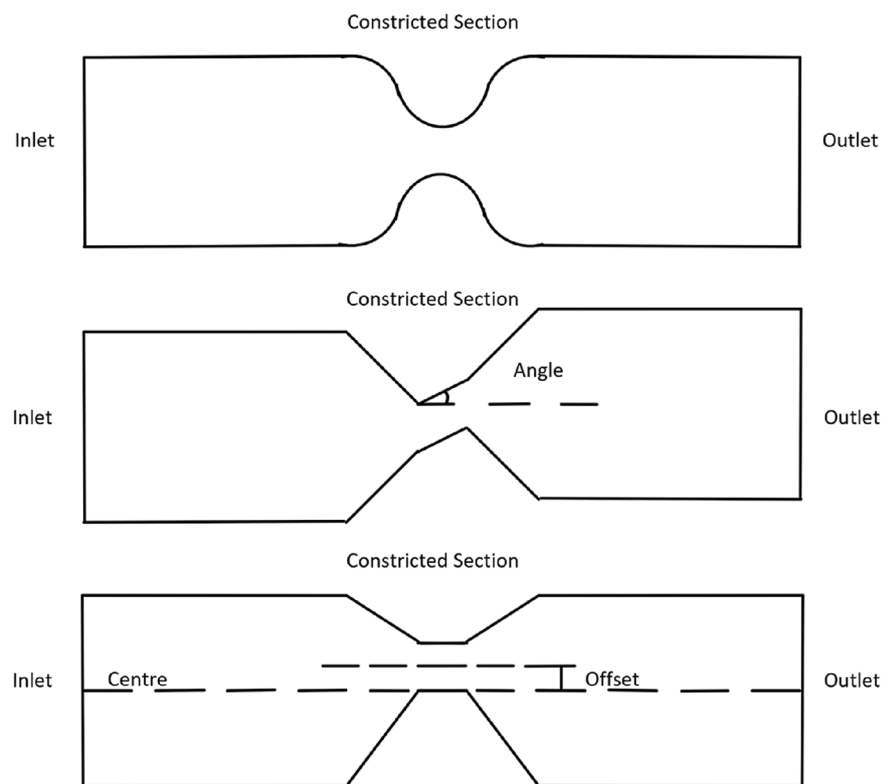

Table 1 Input parameters for basic geometries

No.	v_I	R_I	R_C	v_{peak}
1	0.30	5.0	1.00	3.82
2	0.15	5.0	1.00	1.92
3	0.50	5.0	1.50	2.88
4	0.10	5.0	0.75	2.29
5	0.10	5.0	0.60	3.58
6	0.45	8.0	2.00	2.22
7	0.45	6.0	2.00	2.09
8	0.10	8.0	1.00	2.27
9	0.15	10.0	2.00	2.62
10	0.10	10.0	1.50	2.96

From left to right: the geometry number, the maximum inlet velocity (v_I) in m s^{-1}, the inlet radius (R_I) in mm, the constricted section radius (R_C) in mm, and the peak constricted section velocity (v_{peak}) in m s^{-1}

Table 2 Input parameters for angled/offset geometries

No.	θ	δ	Direction	v_{peak}
11	20.0	0.00	Upward	3.47
12	40.0	0.00	Upward	4.65
13	0.0	1.50	–	3.27
14	0.0	3.00	–	3.18
15	20.0	1.50	Upward	3.37
16	40.0	1.50	Upward	4.32
17	20.0	3.00	Upward	3.28
18	40.0	3.00	Upward	4.12
19	30.0	2.25	Sideways	3.88
20	30.0	2.25	Downward	4.48

From left to right: the geometry number, the angle between the constricted section and the direction normal to the inlet surface (θ) in °, the offset between the constricted section and the centre of the cylinder (δ) in mm, the direction that the constricted section is angled, and the peak constricted section velocity (v_{peak}) in m s^{-1}

present for AR to occur. Figure 1 shows a sketch of the basic geometry.

In total, 20 different simulations were generated using 20 different geometries with a time step size of 0.001 s. The first set of 10 geometries (see Table 1) had variations in inlet velocity, inlet radius, and constricted section radius, while the other set of 10 geometries (see Table 2)

had different shapes. These geometries captured maximum jet velocities between 2.0 and 5.0 m s^{-1} [13, 32]. The second set of geometries were based on the third geometry, as this geometry was reasonably small with minimal computational time. This set had diagonal and off-centre constricted sections, as well combinations of both, to model eccentric jets. Figure 1 demonstrates these differences in shape.

Boundary conditions

The relevant boundary conditions are those relating to the inlet, outlet, and inner wall. The inlet boundary conditions take the most work to define as the velocity varies both in space and time. For blood flowing through the aortic valve, the flow is expected to be fully developed, that is, the velocity is zero at the inner walls due to friction and at its highest in the centre. This can be achieved by extending the length that the blood needs to travel from the inlet to the constricted section, allowing the flow to fully develop before reaching the aortic valve. However, extending the length means the geometry becomes larger, hence increasing the computational time. To compensate for this, the velocity profile at the inlet was defined with a parabolic shape and the upstream length was set to 20 mm [33]. This means the flow will start out more developed than with a uniform profile, and become fully developed before reaching the aortic valve. The velocity was also time-dependent and represented the diastole (where regurgitation occurs), with a rapid initial increase in magnitude before slowly decreasing [13].

The parabolic velocity profile at the inlet in Cartesian coordinates was defined by

$$v = v_I - \frac{x^2 + y^2}{R_I^2 v_I^{-1}}, \tag{1}$$

where v_I is the maximum velocity at the centre of the inlet and R_I is the inlet radius. The remaining boundary conditions were for the outlet and walls. The outlet was defined as an opening with zero pressure difference and the walls were defined as non-permeable with no-slip boundary conditions.

Data preparation

To start, the raw CFD simulations and geometries, which each had 71 timeframes, were sampled onto a uniform Cartesian grid (0.1 mm spacing projected onto multiple planes) to be used as HR images ($392 \times 48 \times 48$ mm). To separate the fluid and non-fluid regions for better data processing and result quantification, binary masks were generated using k-nearest-neighbours [34]. The main difference between these synthetic HR images and MR images relates to the voxel size, VENC, and the amount of noise—HR images are noise-free whereas MR contain phase noise. The LR MR images were obtained from the HR images using the same method as in 4DFlowNet [2] to simulate 4× downsampled MR images with appropriate noise and VENC. This gives the paired LR and HR synthetic images used in network training.

Data augmentation

To augment the data set, similar techniques were used as in 4DFlowNet [2] for each time frame. VENC values were randomly chosen from a set of velocities between 0.3 and 6.0 m s⁻¹, spaced by 0.3 m s⁻¹, for each velocity component. Aliasing was mostly avoided by choosing a VENC larger than the peak velocity. However, since velocity jets cannot be estimated beforehand for actual AR cases (which may cause phase aliasing), a VENC lower than the maximum velocity was chosen with a 10% probability, randomly selecting between 0.3 and 0.6 m s⁻¹ lower. Constant intensity values between 60 and 240 were randomly chosen for the magnitude image, and noise levels were added depending on the SNR, which were randomly and uniformly chosen between 14 and 17 dB.

Since there were a limited number of geometries, further augmentation came in patch generation. From each time frame, 10 patches of $12 \times 12 \times 12$-voxel cubes from the LR image were selected randomly with a minimum fluid region of 20%. These patches acted as random translations, so no extra translation steps were taken. On top of this, for each patch generated another randomly rotated version of the patch was also created. This resulted in 20 patches generated from each time frame and thus 1420 patches per geometry.

Training and validation

To investigate the effect of additional geometries regarding SR image quality, a subset of the data consisting of patches from only five geometries was compared against a the entire data set consisting of patches from all geometries. The validation set in both cases was the same, consisting only of patches from a single geometry.

To investigate the effect of aliasing, duplicates of the two training and validation sets were generated, but with a 10% probability of having a VENC lower than the maximum velocity in any time frame. Networks trained using these aliased data sets were validated against the previously generated validation set without any aliasing, as well as a newly generated validation set with full aliasing, that is, with each time frame having a VENC lower than the maximum velocity.

Network architecture

The simulated pairs of LR and HR synthetic images were used to train a similar deep residual network structure to the one in 4DFlowNet. This consisted of several residual blocks surrounding a central upsampling layer, with the preceding blocks in the LR space pre-processing and acting as denoisers for the input while the following blocks in the HR space refine the output. In 4DFlowNet, LR patches of 16-voxel

cubes were used as input and SR patches of 32-voxel cubes were generated as output, with an upsample factor of 2.

Several changes were made to the above 4DFlowNet architecture to provide higher resolution images with improved accuracy. Firstly, the upsample factor was increased to 4 and the sizes of the input and output patches were changed to 12-voxel and 48-voxel cubes, respectively. The smaller patches account for smaller vessel sizes in the cardiovascular space around the aortic valve [22]. Secondly, the dense and cross stage partial blocks in DenseNet and CSPNet, respectively, were experimented with by using them in place of the 12 residuals blocks in the original 4DFlowNet architecture. The growth rate [30], defined as the number of feature maps in each convolutional layer, of the dense and CSP blocks was set to 16, a quarter of the number of channels in each convolutional layer from the original residual blocks. The adapted 4DFlowNet architecture with residuals blocks (4DFlowNet-Res), with dense blocks (4DFlowNet-Dense), and with cross stage partial blocks (4DFlowNet-CSP) had 3.34, 2.55, and 2.08 million parameters, respectively. These modified networks were implemented with TensorFlow 2.0 [35] and trained using an Adam optimiser [36], with an initial learning rate of 10^{-4} and decay rate of $\sqrt{2}$ after every 14 epochs. Batch sizes of 16 were used, with training completed in 200 epochs.

Loss function

The network was optimised by minimising the mean squared error (MSE) between the paired HR images and the SR images generated from the corresponding input LR ones. The voxel-wise loss was calculated as the mean sum of squared differences between each velocity component:

$$L_{MSE} = \frac{1}{N} \sum_{i=1}^{N} (v'_{x_i} - v_{x_i})^2 + (v'_{y_i} - v_{y_i})^2 + (v'_{z_i} - v_{z_i})^2, \quad (2)$$

where N is the total number of voxels in the geometry, v'_j is the predicted SR velocity, and v_j is the actual HR velocity, for $j \in \{x, y, z\}$.

The MSE of fluid and non-fluid regions were calculated as separate terms due to the imbalance and irregularity of these regions within a specific patch. This gives the total loss to be:

$$L_{total} = L_{MSE_F} + L_{MSE_N}, \quad (3)$$

where L_{MSE_F} and L_{MSE_N} are the voxel-wise loss for the fluid and non-fluid regions, respectively.

The original loss function in 4DFlowNet contained a weighted velocity gradient term to smoothen the gradient between neighbouring vectors [2]. This was omitted from the above loss function as improvements were observed in near-wall velocity estimates with its removal [22].

Evaluation metric

The relative speed error (RE), the relative difference between the SR velocity magnitude (speed) compared to the actual HR speed on the validation set, was used to measure network performance and save model checkpoints. This was only calculated in fluid regions to avoid zero division error, as well as adding a small number ($\epsilon = 10^{-4}$) to the denominator. Furthermore, since many speed values in the HR images were quite small, this could risk significantly over-penalising the model. Thus, an arctangent approach [37] was adopted, giving the following equation for relative speed error:

$$RE = \frac{1}{N} \sum_{i=1}^{N} \arctan\left(\frac{\sqrt{(v'_{x_i} - v_{x_i})^2 + (v'_{y_i} - v_{y_i})^2 + (v'_{z_i} - v_{z_i})^2}}{\sqrt{v_{x_i}^2 + v_{y_i}^2 + v_{z_i}^2} + \epsilon} \right), \quad (4)$$

where N is the total number of voxels in the fluid domain, v'_j and v_j are the predicted SR and actual HR velocities, respectively, for all $j \in \{x, y, z\}$, and 'arctan' is the arctangent function, defined for all real values from negative infinity to infinity with $\lim_{x \to \infty} \tan^{-1} x = \frac{\pi}{2}$ for arctan x.

In addition to the RE, network performance was also evaluated using the root MSE (RMSE) and the structural similarity (SSIM) metric [38] in all three Cartesian velocity components. These were compared against the baseline 4DFlowNet model that had been trained with an upsample factor of 2.

Results

Training was performed using a Tesla V100 GPU with 32 GB memory with networks being trained for 200 epochs. Improvements in relative speed error (RE) plateaued around the 100 epoch mark for 4DFlowNet-Res while still improving for 4DFlowNet-CSP and 4DFlowNet-Dense up till the very last epoch. This can be seen in Fig. 2. The time taken was dependent on the type of network; 4DFlowNet-Res, 4DFlowNet-CSP, and 4DFlowNet-Dense took approximately 163, 168, and 255 h, respectively. Note that these times were for the networks trained using all geometries. For the networks trained using only five geometries, denoted as 4DFlowNet-Res5, 4DFlowNet-CSP5, and 4DFlowNet-Dense5, the times taken were approximately 38, 40, and 63 h, respectively. There was no significant difference in
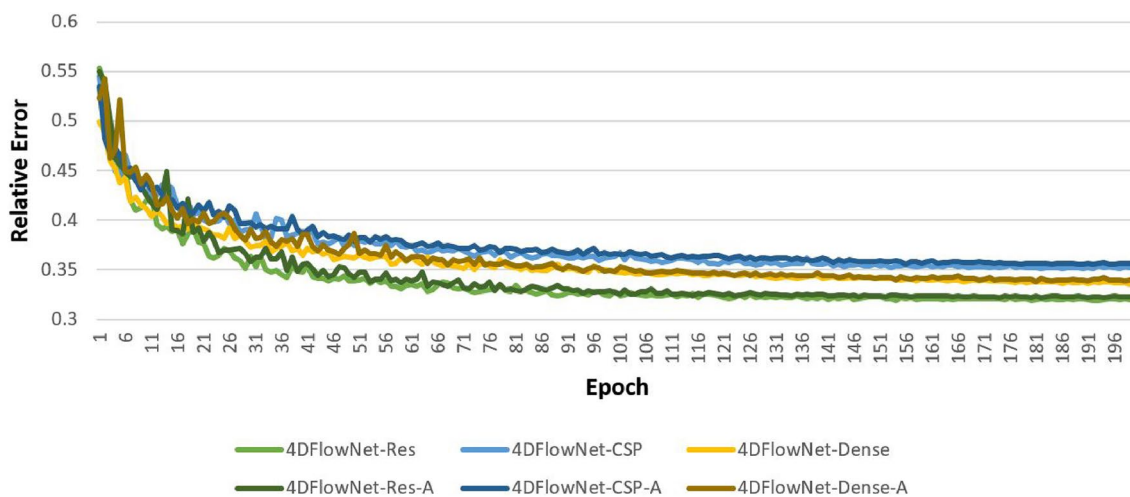
Fig. 2 Relative error across all 200 epochs during training for each network. Networks trained with five geometries were not included as the general trend seen was the same

training time between networks trained with and without a portion of aliased data, denoted by '-A'.

Networks were tested on one complete geometry consisting of 71 timeframes with no (phase/velocity) aliasing, and the same complete geometry with full aliasing. These predictions were required to be patch-based since patches were used as the input and output for each network. The complete geometry was reconstructed by stitching together multiple SR velocity field patches, which was done with a stride of $(n - 4)$ in each Cartesian direction where n is the arbitrary patch size. To avoid patch artifacts at the boundary, four voxels were stripped from each patch side.

The full implementation is accessible on GitHub under an open-source (MIT) license at github.com/dlon450/4DFlowNetv2, with training data available on request.

Synthetic MR images

SR images were analysed visually and quantitatively to better understand how each model was performing. Figure 3 is a visual example of the prediction for the different networks in the constricted section at the peak flow. These display the effectiveness of each network in reducing noise, with the predictions looking quite similar to the ground truth. Furthermore, the networks trained with a proportion of aliased data seem to be performing better than networks without, especially for data with aliasing error.

The values for each evaluation metric were collected and compared in "Appendices 1 and 2" to quantify the performance of every model. Again, these values were taken from the time frame with peak flow, with peak velocities of 2.186, 0.355, and 0.349 m s^{-1} for velocity components v_x, v_y, and v_z, respectively.

The metrics were plotted in "Appendix 3". Briefly, 4DFlowNet-Dense-A and 4DFlowNet-CSP-A seem to perform the best with the lowest RMSE error and largest SSIM in the principle flow direction (v_x), respectively, on both aliased and non-aliased data. However, there does appear to be considerable variation in these results between different networks, depending heavily on the size of the data set and whether aliasing is present.

Regarding the RMSE in the principle flow direction (RMSE$_x$), all networks perform better, on average, than the base 4DFlowNet with less variation in RMSE$_x$. For the smaller data set with five geometries, the residual-based (Res) versions seem to have the smallest RMSE$_x$. There is also noticeable improvement in error between networks trained on the smaller and larger datasets when predicting on non-aliased data. However, when predicting on aliased data, the improvement is significantly more apparent. The increase in dataset size improves the RMSE in the other two flow directions for all networks too, with the CSP versions somewhat better than other versions.

Regarding the SSIM in the principle flow direction (SSIM$_x$), all networks also perform better, on average, than the base 4DFlowNet. However, the SSIM$_x$ seems to worsen for a larger dataset, in general, when predicting on non-aliased data. On the other hand, when predicting on aliased data, there does appear to be slight improvement in SSIM$_x$ across all networks. For the SSIM in the other two flow directions, the values are considerably more varied and worse than in the principle direction. Finally, all networks have a substantially lower RE than the base, with the Dense versions having the lowest RE.

The regression plots in Fig. 4 show that there is exceptional correlation between the SR and synthetic HR images. The regression slopes are very close to one, in the principle

Fig. 3 Predictions on an LR patch from the synthetic 4D flow MRI phase image for different networks with and without aliasing error, focused on the constricted section at two different views along the width (top) and length (bottom) of the geometry. A 2D slice, along the length of the geometry, of the velocity magnitude is shown from the 3D patch for visualisation purposes. Scale is in m s^{-1}
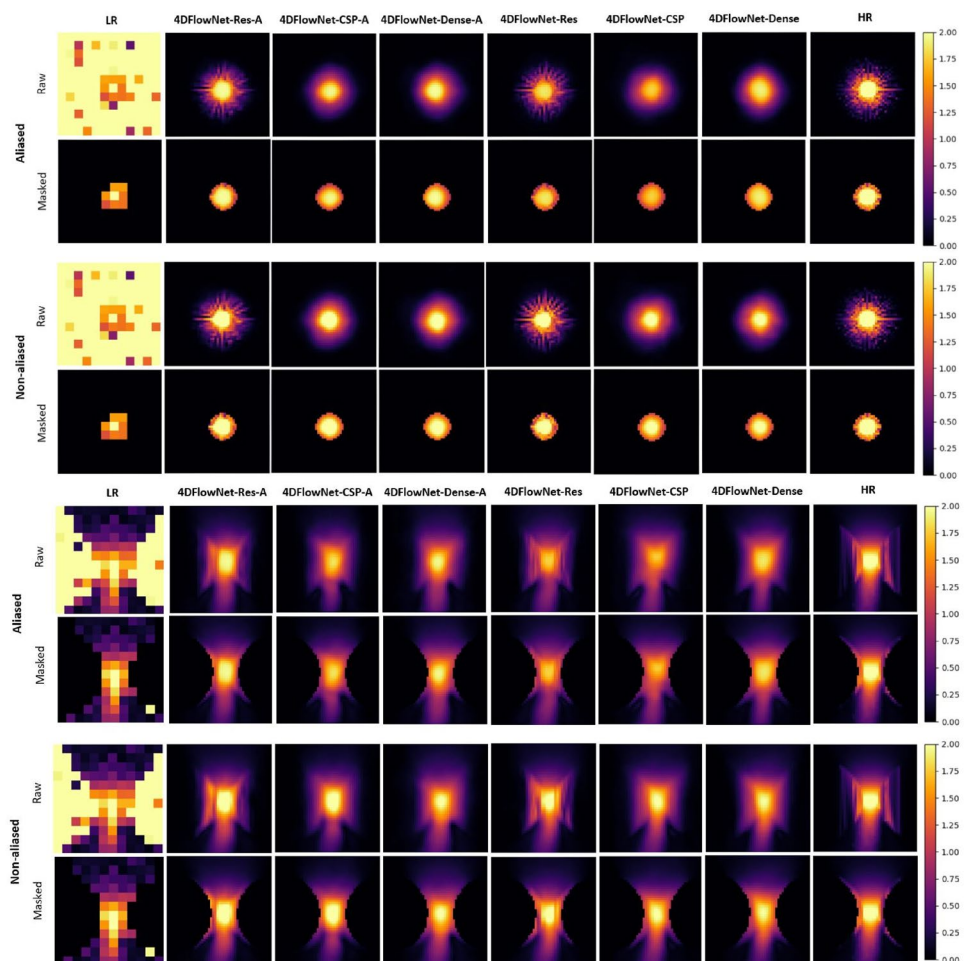
flow direction and velocity magnitude plots, for the two CSP networks. Moreover, offset values are also essentially zero in all examples shown. Training with all non-aliased images appears to have an effect when predicting the higher velocity values in aliased images, with these values being slightly underestimated, as seen on the right in Fig. 4 and confirmed by the Bland–Altman plots too. These plots also indicate minimal bias as the deviations appear constant, uniform, and are all less than 0.08 m s^{-1}.

For velocities within the constricted section, Fig. 5 shows the correlation between SR and synthetic HR images. Although a slight underestimation bias seem to prevail, 4DFlowNet-CSP shows noticeable improvement from the baseline 4DFlowNet improving on the otherwise observed deviations at higher velocities. Note that the general trends and comments regarding these regression and Bland–Altman plots were present in all other evaluated networks too.

In vivo 4D flow MRI data

Ethical approval for this study was granted by the Health and Disability Ethics Committee of New Zealand (17/CEN/226), and written informed consent was obtained

from each participant. Two sets of in vivo 4D flow MRI data (CH34 and CH37) were acquired and used exclusively to show how the network would perform, seen in Fig. 6. CH34 (resp. CH37) had an AR VTi of 264.8 cm (resp. 201.2 cm), VC width of 3 mm (resp. 3 mm), peak velocity of 4.5 m s^{-1} (resp. 4.2 m s^{-1}), and pressure drop of 81 mmHg (resp. 70.5 mmHg).

These datasets were LR images only, with no HR images available to compare the predicted SR images against. However, the SR image produced was compared against the original 4D flow MRI image to help better understand the prediction. The SR image seems to have effectively removed noise from the LR image and strengthened the AR signal. Other networks performed similarly, with predicted SR images almost identical to the one shown. The top set of data (CH34) was also segmented and visualised with Para-View [39], shown in Fig. 7. Image stitching appears to have been performed correctly, with velocities within the fastest section preserved well.
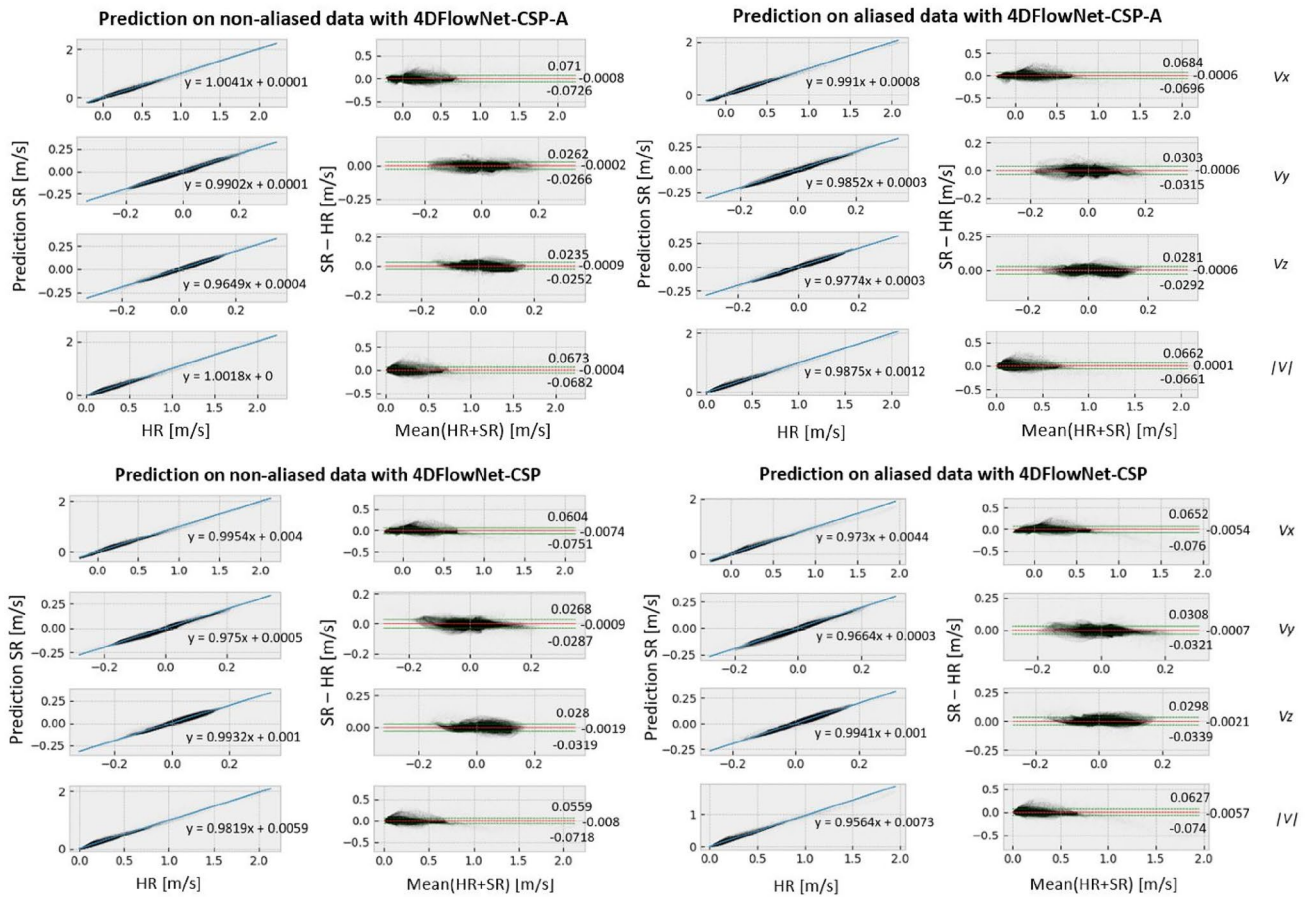
Fig. 4 Regression and Bland–Altman plots for prediction on non-aliased (left) and aliased (right) data with 4DFlowNet-CSP-A (top) and 4DFlowNet-CSP (bottom). These plots are for each of the velocity components and magnitude (v_x, v_y, v_z, and $\|v\|$, respectively, from top to bottom) between SR and synthetic HR images
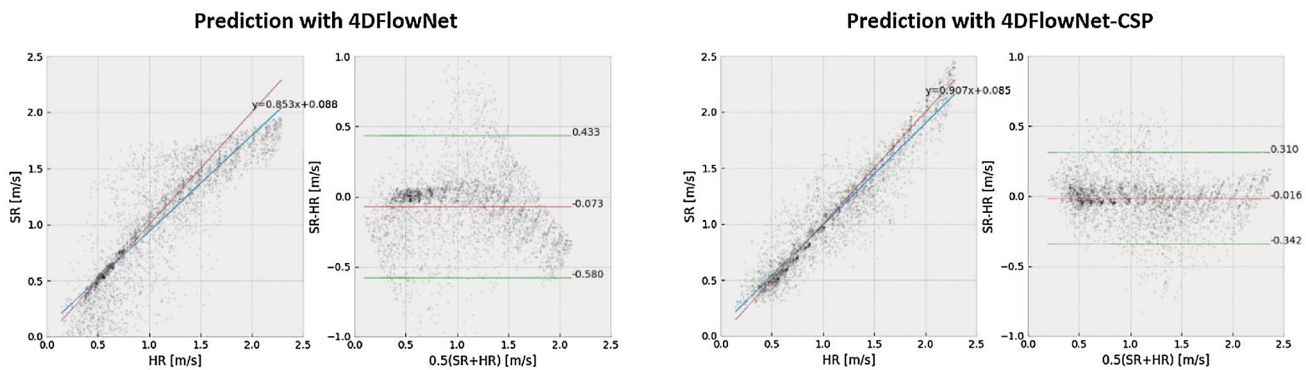

Fig. 5 Regression and Bland–Altman plots for $\|v\|$ prediction within the constricted section, comparing the baseline 4DFlowNet (left) against the adapted 4DFlowNet-CSP (right)

Discussion

A noteworthy consideration when analysing the results is that the focus should be more towards metrics and values obtained in the principle flow direction x. Since the peak velocities in the other flow directions, y and z, were almost 10 times smaller than that in the principle direction, networks would have difficulty differentiating between velocity and noise. This is evident in the y and z RMSE values, which were over half of their corresponding peak velocities. The y

Fig. 6 Predicted SR images on actual 4D flow MRI data. Two sets of 4D flow MRI data were used (top and bottom), with the original LR images on the left and the SR images in the other three columns, titled by the network used for prediction. These are 2D slices of the 3D image, showing the velocity magnitude taken at peak flow. Scale is in m s^{-1}
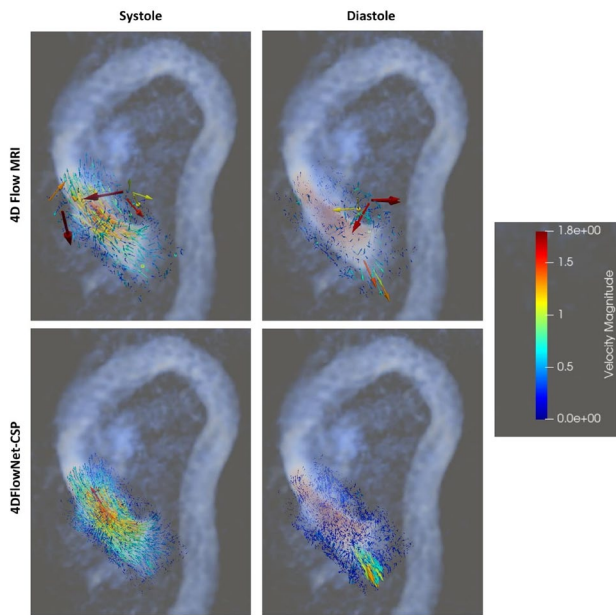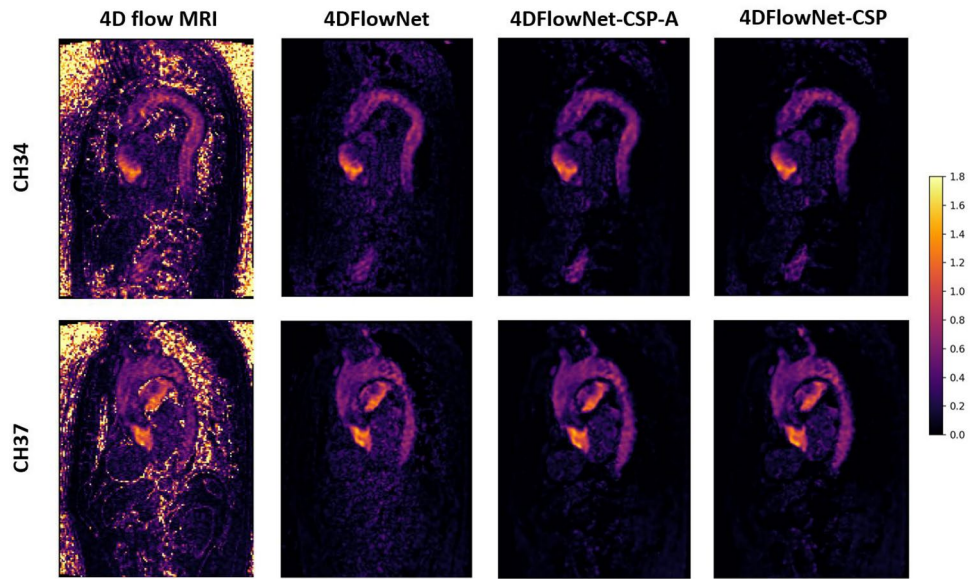

Fig. 7 Original 4D flow MRI image against the SR image produced by 4DFlowNet-CSP at systole and diastole, visualised in ParaView. The regurgitant jet can be clearly seen during diastole in the SR image produced. Scale is in m s^{-1}

and z SSIM values were also much worse, potentially exhibiting strange behaviour in these low velocity fields too [40].

Synthetic 4D flow MRI data

The main limitation of previous work has been related to insufficient data and flow characteristics [2, 20, 22]. To understand the effect of additional flow characteristics in

the dataset, a wider range was used in this study. There was definitely noticeable improvement in the RMSE across all networks tested, particularly when predicting on aliased data. This indicates that with more geometries and hence flow characteristics, the network seems to generalise better and is more robust against data it has not seen.

In terms of the synthetic LR and HR image pairs, the downsampling process and patch-based approach seemed to work effectively. The noisy LR patches enabled the network to learn noise removal while also enabling greater generalisation to unknown flow characteristics or geometries [2]. However, a small portion of these incorrect non-fluid velocity areas around the fluid domain appear to have been included in network training, seen in Fig. 3. These were generated during the linear interpolation process when obtaining the HR images from the CFD data. This suggests that the binary mask for separating between the fluid and non-fluid regions, created using k-nearest-neighbours, needs improvement. An obvious consequence is that the network may learn incorrect flow characteristics and predict less accurately.

Incorporating aliased patches into the training data seemed to work effectively as well, improving the RMSE across networks trained with all geometries. This was done by choosing VENC values lower than the maximum velocity, within a particular time frame, with a 10% probability. Note that this probability was chosen arbitrarily. For the VENC, if it is set too high, visualization of the jet may not be obtained and be inaccurate, as well as having poorer SNR. On the other hand, if it is set too low, flow characteristics may be lost and a mosaic pattern will be shown [41]. This means that for the time frames with aliasing, the velocities lower than the VENC would have been captured

significantly more clearly, with only a small portion of high velocity characteristics being lost. For these time frames, the purpose would be to help the network better learn the flow characteristics in lower velocity fields, improving robustness and generalisability. However, this hypothesis has not been proven yet and more experiments on aliasing will be required to fully understand its effect on network performance. This would involve varying its probability of occurring as well as choosing VENC values considerably lower than the maximum velocity. Future work may also include a thorough analysis of patch size to quantify its effect on network accuracy.

Network architecture

A major drawback of residual blocks is that they suffer from limited learning ability [30]. This was seen in the results, in which the residual networks did not show much improvement even after adding many more geometries or introducing aliased images into the training data. Furthermore, the RE for residual networks seemed to plateau at around the 100th epoch, whereas the RE for the other networks seemed to continue improving, although at a decreasing rate, even up till the last epoch. These observations can be seen in Fig. 2. Despite this, the residual networks still performed well, with RMSE, SSIM, and RE values similar to the other networks, as well as having the fastest training and prediction times. This suggests that, when data is insufficient or very limited, residual networks may work best.

With sufficient or abundant data, cross stage partial or dense networks may be preferred. The learning ability in these types of networks have a significantly higher ceiling than residual networks [30, 31], with the only difference between these two network structures being the training and prediction times. Cross stage partial networks had training and prediction times almost as fast as residual networks, whereas dense networks were almost 1.6 times slower. Although training times may not be a crucial problem, clinicians and patients alike may require and desire fast prediction times. This leads to cross stage partial networks being preferred. Otherwise, if training and prediction times are not significant constraints, then dense networks may be the optimal choice. Furthermore, the growth rate for the dense and cross stage partial networks was set quite low, at a quarter of the number of feature maps in each convolutional layer within the residual blocks. This limits the learning ability of these networks, as there are significantly less parameters, so testing larger values of this hyperparameter will be beneficial and likely improve network performance. Similarly, only a quarter of feature maps were taken from the base input

layer within each partial dense block within the cross stage partial networks, so larger values of this hyperparameter will likely improve performance as well.

Despite the network architecture seeming to work quite effectively, there are still improvements that could be made on top of modifying the residual blocks. Presently, the network is not taking full advantage of the temporal aspect in 4D flow MRI. Modifying it to incorporate characteristics of recurrent neural networks [42] may help the network understand this temporal aspect better. This could be done by using predictions of the same patches from one or two time frames prior. Additionally, including physical properties of fluids may also help the network in learning flow characteristics, as seen in [21, 22]. However, due to the bulky nature of the velocity data, which were 3D volumes for each velocity component, these ideas were not considered further as it would have been too costly to process given the available resources.

Clinical application

The clinical motivation for the current study was the assessment of regurgitant or highly stenotic valvular flow; instances where quantification of regional velocities and changes in pressure act as effective biomarkers to describe the severity of disease. In these instances, and in clinical practice, assessment of *peak* velocities through the vena contracta are used to symbolize disease severity. In fact, derivation of regional pressure drops are routinely derived from such peak velocity measures using the so called simplified Bernoulli equation [43] (coupling peak velocities to effective pressure changes). In this light, our results bear possible clinical implications in that SR velocities effectively recovers HR reference measures. This holds true throughout the evaluation domain, and in the clinically important constricted section, underestimation biases associated with the original 4DFlowNet formulation were effectively suppressed as seen in Fig. 5. The effect of the remaining minor deviation from a true 1:1 correlation between HR and SR data remains to be explored in a larger clinical setting, and in more complex flow scenarios—such as in the instance of regurgitant flow—higher-order methods might be required to derive pressure drops from measured velocity data [4, 44]. Nevertheless, just as we have shown the potential of recovering functional hemodynamic behavior through the spatially challenging cerebrovascular space using SR 4D flow MRI [22], the results of the current study bear similar potential in recovering clinically relevant hemodynamic metrics in aortic regurgitation.

Limitations

Although the number of geometries and flow patterns have significantly increased from previous studies, a wider range of characteristics can still be included. Additional data will diversify the data set further and improve the model's robustness and generalisability even more. On top of this, more testing on real 4D flow MRI data will be required to validate model performance. Currently, this validation is only done by visually analysing model predictions to see if they look reasonable and sensible. The preferred approach would to be validate quantitatively with a pair of corresponding LR and HR 4D flow MRI images, calculating metrics such as RMSE and SSIM to properly understand model performance. Additional metrics may also be required to properly diagnose aortic regurgitation severity such as the effective regurgitant orifice area, which would be measured as part of a clinical workflow. Furthermore, blood flowing through the aortic valve was expected to be fully developed, which may not be the case in reality.

Lastly, several practical limitations can also be noted. Due to the limited GPU resources, training took approximately 60 days for all networks. With either more time or increased GPU resources, more networks could be trained by using additional optimizers, adding a weighted term to the loss function, as well as different values for hyperparameters such as the growth rate, the probability that aliasing occurs, and the upsample factor; a sensitivity test for this may be worthwhile. Moreover, networks could also be trained for more epochs, or until the error plateaus, to better gauge the learning ability of different networks. Finally, memory constraints were a factor too, as an upsample factor of 4 led to 64 times more usage in disk space. This would not be feasible for much larger upsample factors, so a different data representation may be required in future work.

Conclusion

In this study, 4DFlowNet was enhanced and adapted to effectively quantify hemodynamic metrics for AR by producing super-resolution 4D flow MRI images with an upsample factor of 4. The results show that by adding more geometries and hence flow characteristics into the data set, the accuracy of 4DFlowNet predictions are improved. Moreover, the comparison of different network architecture suggested that the original residual network structure limits learning ability and can be further refined.

Appendix 1

See Table 3.

Table 3 Summary of prediction errors and evaluation metrics for different networks, when predicting on non-aliased data

Network	$RMSE_x^a$	$RMSE_y^a$	$RMSE_z^a$	$SSIM^b$	RE^c
4DFlowNet	0.0574 ± 0.0574	0.0198 ± 0.0198	0.0177 ± 0.0177	(0.832, 0.708, 0.705)	0.543
4DFlowNet-Res	0.0361 ± 0.0360	0.0138 ± 0.0138	0.0135 ± 0.0135	(0.923, **0.843**, 0.820)	0.330
4DFlowNet-CSP	0.0354 ± 0.0346	0.0142 ± 0.0142	0.0154 ± 0.0153	(0.915, 0.728, 0.766)	0.316
4DFlowNet-Dense	0.0372 ± 0.0359	0.0138 ± 0.0138	0.0139 ± 0.0138	(0.908, 0.794, 0.789)	**0.296**
4DFlowNet-Res-A	0.0369 ± 0.0363	0.0136 ± 0.0136	0.0138 ± 0.0136	(0.924, 0.829, 0.824)	0.340
4DFlowNet-CSP-A	0.0366 ± 0.0366	$\mathbf{0.0135 \pm 0.0135}$	$\mathbf{0.0126 \pm 0.0124}$	(**0.939**, 0.812, 0.803)	0.330
4DFlowNet-Dense-A	$\mathbf{0.0350 \pm 0.0343}$	0.0149 ± 0.0148	0.0163 ± 0.0163	(0.912, 0.747, 0.705)	0.301
4DFlowNet-Res5	0.0375 ± 0.0375	0.0160 ± 0.0158	0.0163 ± 0.0162	(0.934, 0.796, *0.825*)	*0.344*
4DFlowNet-CSP5	0.0399 ± 0.0397	0.0152 ± 0.0152	0.0145 ± 0.0142	(*0.935*, 0.812, 0.804)	0.338
4DFlowNet-Dense5	0.0407 ± 0.0406	0.0166 ± 0.0166	0.0150 ± 0.0147	(0.929, 0.764, 0.799)	0.352
4DFlowNet-Res5-A	0.0415 ± 0.0415	0.0196 ± 0.0196	0.0189 ± 0.0189	(0.924, *0.815*, 0.817)	0.354
4DFlowNet-CSP5-A	0.0438 ± 0.0438	0.0219 ± 0.0219	0.0211 ± 0.0211	(0.929, 0.803, 0.811)	0.362
4DFlowNet-Dense5-A	0.0402 ± 0.0402	0.0194 ± 0.0193	0.0169 ± 0.0169	(0.928, 0.807, 0.784)	0.364

The best values for each metric are in bold. The best values for networks trained with only five geometries have also been italic

[a]$RMSE_i \pm$ s.d. for velocity component v_i, in m s^{-1}

[b]Includes all three Cartesian velocity components in the form (v_x, v_y, v_z)

[c]Or equivalently, MAAPE

Appendix 2

See Table 4.

Table 4 Summary of prediction errors and evaluation metrics for different networks, when predicting on aliased data

Network	RMSE$_x^1$	RMSE$_y^1$	RMSE$_z^1$	SSIM^2	RE^3
4DFlowNet	0.0642 ± 0.0641	0.0292 ± 0.0292	0.0272 ± 0.0272	(0.855, 0.724, 0.814)	0.512
4DFlowNet-Res	0.0361 ± 0.0361	**0.0142 ± 0.0142**	**0.0138 ± 0.0138**	(0.927, 0.858, 0.836)	0.334
4DFlowNet-CSP	0.0364 ± 0.0360	0.0161 ± 0.0161	0.0164 ± 0.0163	(0.917, 0.795, 0.778)	0.309
4DFlowNet-Dense	0.0356 ± 0.0344	0.0153 ± 0.0153	0.0159 ± 0.0159	(0.917, 0.812, 0.798)	0.309
4DFlowNet-Res-A	0.0355 ± 0.0350	0.0148 ± 0.0148	0.0140 ± 0.0138	(0.925, 0.845, 0.848)	0.342
4DFlowNet-CSP-A	0.0352 ± 0.0352	0.0158 ± 0.0158	0.0147 ± 0.0146	(**0.942**, 0.845, **0.849**)	0.313
4DFlowNet-Dense-A	**0.0340 ± 0.0343**	0.0164 ± 0.0148	0.0164 ± 0.0163	(0.922, 0.785, 0.757)	**0.284**
4DFlowNet-Res5	0.0468 ± 0.0468	0.0176 ± 0.0176	0.0170 ± 0.0169	(0.865, 0.899, 0.582)	*0.337*
4DFlowNet-CSP5	0.0529 ± 0.0529	*0.01743 ± 0.01743*	0.01636 ± 0.01619	(0.912, 0.921, 0.509)	0.341
4DFlowNet-Dense5	0.0508 ± 0.0508	0.0178 ± 0.0177	0.0171 ± 0.0171	(0.902, 0.910, 0.587)	0.353
4DFlowNet-Res5-A	*0.0437 ± 0.0437*	0.0175 ± 0.0175	*0.0159 ± 0.0159*	(0.901, 0.902, 0.522)	0.344
4DFlowNet-CSP5-A	0.0466 ± 0.0465	0.0180 ± 0.0180	0.0164 ± 0.0164	(*0.914*, ***0.925***, *0.610*)	0.350
4DFlowNet-Dense5-A	0.0474 ± 0.0473	0.0194 ± 0.0193	0.0162 ± 0.0162	(0.895, 0.902, 0.506)	0.338

The best values for each metric are in bold. The best values for networks trained with only five geometries have also italic
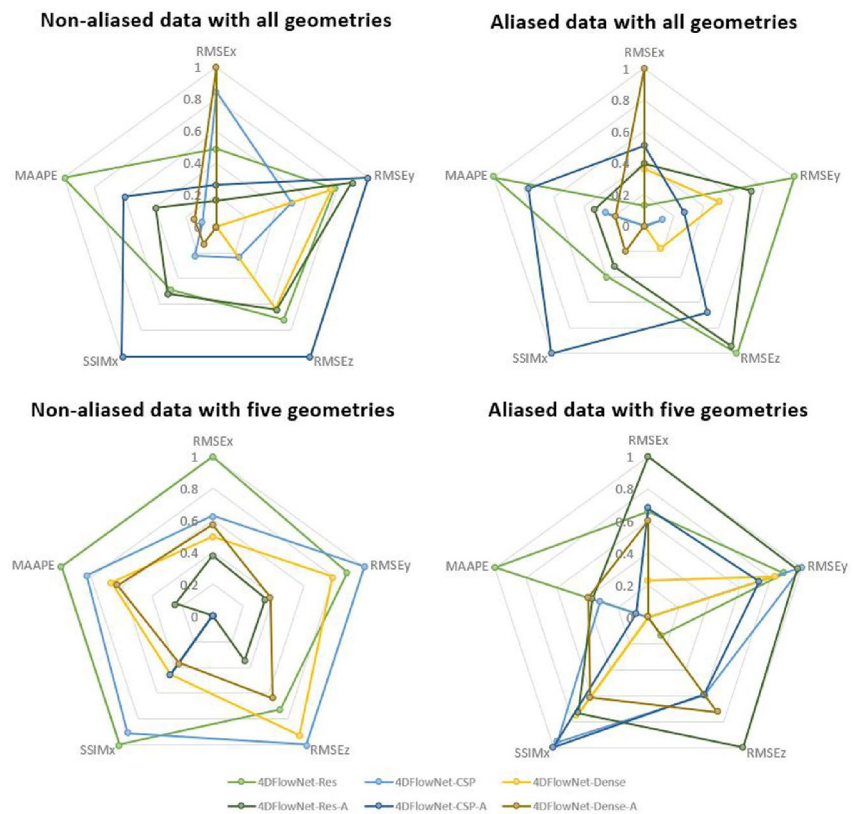
[1]RMSE$_i$ ± s.d. for velocity component v_i, in m s^{-1}

[2]Includes all three Cartesian velocity components in the form (v_x, v_y, v_z)

[3]Or equivalently, MAAPE

Appendix 3

See Fig. 8.

Fig. 8 Comparison of networks for different metrics and results, varied by the data set size and whether aliasing was used. Values were normalised to be within 0 and 1, representing the worst and best values, respectively, for a specific metric relative to other networks

Declarations

References

1. Bekeredjian R, Grayburn PA (2005) Valvular heart disease. Circulation 112(1):125–134. https://doi.org/10.1161/CIRCULATIONAHA.104.488825
2. Ferdian E, Suinesiaputra A, Dubowitz DJ, Zhao D, Wang A, Cowan B, Young AA (2020) 4DFlowNet: super-resolution 4D flow MRI using deep learning and computational fluid dynamics. Front Phys 8:138. https://doi.org/10.3389/fphy.2020.00138
3. Chaliki HP, Mohty D, Avierinos J-F, Scott CG, Schaff HV, Tajik AJ, Enriquez-Sarano M (2002) Outcomes after aortic valve replacement in patients with severe aortic regurgitation and markedly reduced left ventricular function. Circulation 106(21):2687–2693. https://doi.org/10.1161/01.CIR.0000038498.59829.38
4. Marlevi D, Ruijsink B, Balmus M, Dillon-Murphy D, Fovargue D, Pushparajah K, Bertoglio C, Colarieti-Tosti M, Larsson M, Lamata P, Figueroa CA, Razavi R, Nordsletten DA (2019) Estimation of cardiovascular relative pressure using virtual work-energy. Sci Rep 9(1):1375. https://doi.org/10.1038/s41598-018-37714-0
5. Markl M, Frydrychowicz A, Kozerke S, Hope M, Wieben O (2012) 4D flow MRI. J Magn Reson Imaging 36(5):1015–1036. https://doi.org/10.1002/jmri.23632
6. Itatani K, Miyazaki S, Furusawa T, Numata S, Yamazaki S, Morimoto K, Makino R, Morichi H, Nishino T, Yaku H (2017) New imaging tools in cardiovascular medicine: computational fluid dynamics and 4D flow MRI. Gen Thorac Cardiovasc Surg 65:611–621
7. Alvarez A, Martinez V, Pizarro G, Recio M, Cabrera JÁ (2020) Clinical use of 4D flow MRI for quantification of aortic regurgitation. Open Heart. https://doi.org/10.1136/openhrt-2019-001158
8. Catapano F, Pambianchi G, Cundari G, Rebelo J, Cilia F, Carbone I, Catalano C, Francone M, Galea N (2020) 4D flow imaging of the thoracic aorta: is there an added clinical value? Cardiovasc Diagn Ther. https://doi.org/10.21037/cdt-20-452
9. Bohbot Y, Renard C, Manrique A, Levy F, Maréchaux S, Gerber BL, Tribouilloy C (2020) Usefulness of cardiac magnetic resonance imaging in aortic stenosis. Circ Cardiovasc Imaging 13(5):010356. https://doi.org/10.1161/CIRCIMAGING.119.010356
10. Minderhoud SCS, Wentzel JJ, Roos-Hesselink JW, Hirsch A (2021) Six-year serial follow-up with aortic 4D flow cardiovascular magnetic resonance in a patient with bicuspid aortic valve. Eur Heart J Cardiovasc Imaging 22(9):144–144. https://doi.org/10.1093/ehjci/jeab058
11. Jones L, Pressdee D, Lamont P, Baird R, Murphy K (1998) A phase contrast (PC) rephase/dephase sequence of magnetic resonance angiography (MRA): a new technique for imaging distal run-off in the pre-operative evaluation of peripheral vascular disease. Clin Radiol 53(5):333–337. https://doi.org/10.1016/s0009-9260(98)80003-2
12. Jiang J, Kokeny P, Ying W, Magnano C, Zivadinov R, Mark Haacke E (2015) Quantifying errors in flow measurement using phase contrast magnetic resonance imaging: comparison of several boundary detection methods. Magn Reson Imaging 33(2):185–193. https://doi.org/10.1016/j.mri.2014.10.009
13. Sallach SM, Reimold SC (2007). In: Solomon SD, Bulwer B (eds) Echocardiographic evaluation of aortic regurgitation. Humana Press, Totowa, pp 223–238. https://doi.org/10.1007/978-1-59259-977-6_12
14. LeCun Y, Bengio Y, Hinton G (2015) Deep learning. Nature 521:436–444. https://doi.org/10.1038/nature14539
15. Maier A, Syben C, Lasser T, Riess C (2019) A gentle introduction to deep learning in medical image processing. Z Med Phys 29(2, Special Issue: Deep Learning in Medical Physics):86–101. https://doi.org/10.1016/j.zemedi.2018.12.003
16. Park SC, Park MK, Kang MG (2003) Super-resolution image reconstruction: a technical overview. IEEE Signal Process Mag 20(3):21–36. https://doi.org/10.1109/MSP.2003.1203207
17. Chen Y, Christodoulou AG, Zhou Z, Shi F, Xie Y, Li D (2020) MRI super-resolution with GAN and 3D multi-level DenseNet: smaller, faster, and better. Single image super-resolution for MRI with deep learning
18. Chen Y, Xie Y, Zhou Z, Shi F, Christodoulou AG, Li D (2018) Brain MRI super resolution using 3D deep densely connected neural networks. In: 2018 IEEE 15th international symposium on biomedical imaging (ISBI 2018), 2018. https://doi.org/10.1109/isbi.2018.8363679
19. Chun J, Zhang H, Gach HM, Olberg S, Mazur T, Green O, Kim T, Kim H, Kim JS, Mutic S, Park JC (2019) MRI super-resolution reconstruction for MRI-guided adaptive radiotherapy using cascaded deep learning: in the presence of limited training data and unknown translation model. Med Phys 46(9):4148–4164. https://doi.org/10.1002/mp.13717
20. Rutkowski DR, Roldán-Alzate A, Johnson KM (2021) Enhancement of cerebrovascular 4D flow MRI velocity fields using machine learning and computational fluid dynamics simulation data. Nature. https://doi.org/10.1038/s41598-021-89636-z
21. Kissas G, Yang Y, Hwuang E, Witschey WR, Detre JA, Perdikaris P (2020) 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 358:112623. https://doi.org/10.1016/j.cma.2019.112623
22. Ferdian E, Marlevi D, Schollenberger J, Aristova M, Edelman ER, Schnell S, Figueroa CA, Nordsletten DA, Young AA (2021) Cerebrovascular super-resolution 4D Flow MRI—using deep learning to non-invasively quantify velocity, flow, and relative pressure. bioRxiv. https://doi.org/10.1101/2021.08.25.457611
23. Zhao C, Shao M, Carass A, Li H, Dewey BE, Ellingsen LM, Woo J, Guttman MA, Blitz AM, Stone M, Calabresi PA, Halperin H,

Prince JL (2019) Applications of a deep learning method for anti-aliasing and super-resolution in MRI. Magn Reson Imaging 64:132–141. https://doi.org/10.1016/j.mri.2019.05.038

24. Lopez-Mattei JC, Shah DJ (2013) The role of cardiac magnetic resonance in valvular heart disease. Methodist Debakey Cardiovasc J 9(3):142–148. https://doi.org/10.14797/mdcj-9-3-142

25. Jiao L, Zhang F, Liu F, Yang S, Li L, Feng Z, Qu R (2019) A survey of deep learning-based object detection. IEEE Access 7:128837–128868. https://doi.org/10.1109/ACCESS.2019.2939201

26. Zhao C, Carass A, Dewey BE, Woo J, Oh J, Calabresi PA, Reich DS, Sati P, Pham DL, Prince JL (2018) A deep learning based anti-aliasing self super-resolution algorithm for MRI. In: Frangi AF, Schnabel JA, Davatzikos C, Alberola-López C, Fichtinger G (eds) Medical image computing and computer assisted intervention—MICCAI 2018, 2018. Springer, Cham, pp 100–108

27. Zhao C, Dewey BE, Pham DL, Calabresi PA, Reich DS, Prince JL (2021) SMORE: a self-supervised anti-aliasing and super-resolution algorithm for MRI using deep learning. IEEE Trans Med Imaging 40(3):805–817. https://doi.org/10.1109/TMI.2020.3037187

28. Pham C-H, Ducournau A, Fablet R, Rousseau F (2017) Brain MRI super-resolution using deep 3D convolutional networks. In: 2017 IEEE 14th international symposium on biomedical imaging (ISBI 2017), 2017, pp 197–200. https://doi.org/10.1109/ISBI.2017.7950500

29. He K, Zhang X, Ren S, Sun J (2015) Deep residual learning for image recognition. CoRR abs/1512.03385. arXiv: 1512.03385

30. Huang G, Liu Z, Weinberger KQ (2016) Densely connected convolutional networks. CoRR abs/1608.06993. arXiv: 1608.06993

31. Wang C, Liao HM, Yeh I, Wu Y, Chen P, Hsieh J (2019) CSPNet: a new backbone that can enhance learning capability of CNN. CoRR abs/1911.11929. arXiv: 1911.11929

32. Tei C, Kisanuki A, Arikawa K, Otsuji Y, Arima S, Kashima T, Tanaka H, Kushitani M, Shah PM (1985) Continuous wave Doppler echocardiographic assessment of aortic stenosis and aortic regurgitation. J Cardiogr 15(1):111–121

33. Madhavan S, Kemmerling EC (2016) The effect of inlet boundary conditions in image-based CFD modeling of aortic flow. In: APS Division of Fluid Dynamics meeting abstracts, 2016. APS meeting abstracts, p 15-001

34. Guo G, Wang H, Bell D, Bi Y, Greer K (2003) KNN model-based approach in classification. In: Meersman R, Tari Z, Schmidt DC (eds) On the move to meaningful Internet systems 2003: CoopIS, DOA, and ODBASE. Springer, Berlin, pp 986–996

35. Abadi M, Barham P, Chen J, Chen Z, Davis A, Dean J, Devin M, Ghemawat S, Irving G, Isard M, Kudlur M, Levenberg J, Monga R, Moore S, Murray DG, Steiner B, Tucker P, Vasudevan V, Warden P, Wicke M, Yu Y, Zheng X (2016) TensorFlow: a system for large-scale machine learning. In: 12th USENIX symposium on operating systems design and implementation (OSDI 16), 2016. USENIX Association, Savannah, pp 265–283

36. Kingma DP, Ba J (2017) Adam: a method for stochastic optimization

37. Kim S, Kim H (2016) A new metric of absolute percentage error for intermittent demand forecasts. Int J Forecast 32(3):669–679. https://doi.org/10.1016/j.ijforecast.2015.12.003

38. Wang Z, Bovik AC, Sheikh HR, Simoncelli EP (2004) Image quality assessment: from error visibility to structural similarity. IEEE Trans Image Process 13(4):600–612

39. Ahrens J, Geveci B, Law C (2005) ParaView: an end-user tool for large data visualization. In: The visualization handbook 717(8). Elsevier, Amsterdam

40. Pambrun J-F, Noumeir R (2015) Limitations of the SSIM quality metric in the context of diagnostic imaging. In: 2015 IEEE international conference on image processing (ICIP), 2015, pp 2960–2963. https://doi.org/10.1109/ICIP.2015.7351345

41. Rebergen SA, van der Wall EE, Doornbos J, de Roos A (1993) Magnetic resonance measurement of velocity and flow: technique, validation, and cardiovascular applications. Am Heart J 126(6):1439–1456. https://doi.org/10.1016/0002-8703(93)90544-J

42. Sherstinsky A (2020) Fundamentals of recurrent neural network (RNN) and long short-term memory (LSTM) network. Physica D 404:132306. https://doi.org/10.1016/j.physd.2019.132306

43. Stamm RB, Martin RP (1983) Quantification of pressure gradients across stenotic valves by Doppler ultrasound. J Am Coll Cardiol 2(4):707–718

44. Marlevi D, Schollenberger J, Aristova M, Ferdian E, Ma Y, Young AA, Edelman ER, Schnell S, Figueroa CA, Nordsletten DA (2021) Noninvasive quantification of cerebrovascular pressure changes using 4D flow MRI. Magn Reson Med 86(6):3096–3110

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.