



Deep variational network for rapid 4D flow MRI reconstruction

Valery Vishnevskiy^{1,2}✉, Jonas Walheim^{1,2} and Sebastian Kozerke¹

Phase-contrast magnetic resonance imaging (MRI) provides time-resolved quantification of blood flow dynamics that can aid clinical diagnosis. Long in vivo scan times due to repeated three-dimensional (3D) volume sampling over cardiac phases and breathing cycles necessitate accelerated imaging techniques that leverage data correlations. Standard compressed sensing reconstruction methods require tuning of hyperparameters and are computationally expensive, which diminishes the potential reduction of examination times. We propose an efficient model-based deep neural reconstruction network and evaluate its performance on clinical aortic flow data. The network is shown to reconstruct undersampled 4D flow MRI data in under a minute on standard consumer hardware. Remarkably, the relatively low amounts of tunable parameters allowed the network to be trained on images from 11 reference scans while generalizing well to retrospective and prospective undersampled data for various acceleration factors and anatomies.

Four-dimensional (4D) flow magnetic resonance imaging (MRI) provides spatiotemporally resolved quantification of blood flow and offers great potential for the assessment of cardiovascular disease, for example aortic valve stenosis, atherosclerosis or vessel wall remodelling¹. However, clinical adaptation of the method has been hampered by the long examination times.

Many efforts have been dedicated to accelerate flow acquisition by exploiting redundancies in the data. Partial Fourier imaging² has been used for moderate acceleration³, but the underlying assumption of a slowly varying phase has been shown to be incorrect for 4D flow MRI⁴. Parallel imaging (PI)⁵, which exploits the spatially varying sensitivity of receiver elements in the coil array, has become a standard for accelerated imaging, but undersampling rates are limited by noise amplification⁶. The advent of compressed sensing (CS)⁷ has enabled acceleration of 4D flow MRI by acquiring only a subset of k -space data and exploiting prior information about data regularities during reconstruction^{8–13}, with typical acceleration factors ranging from 5 (ref. ¹⁰) to 27 (ref. ¹³). In particular, locally low rank (LLR) regularized reconstruction¹⁴ has been successful, iteratively balancing the data fidelity cost and the singular norm of a patch matrix stacked over cardiac phases (see Methods for details). However, iterative reconstruction methods, as used in CS, increase the reconstruction times considerably, implying that evaluation of 4D flow MRI data will typically happen when the subject has already been moved out of the scanner.

In recent years, deep neural networks have gained increasing popularity in MR image reconstruction. In the training stage, the neural network learns abstract features from a set of scans. After training, newly acquired data are reconstructed with very little computational effort by inference with the learned weights. This reduction in reconstruction times can facilitate the use of accelerated imaging methods in clinical practice. Moreover, reconstruction results can be superior to traditional CS methods^{15,16}. Some approaches discard concepts of iterative image reconstruction altogether, for example, by learning end-to-end mappings from k -space to image space¹⁷. As a downside, such networks usually require abundant amounts of high-quality training data, which are not available

for high-dimensional flow MRI. Model-based neural reconstruction networks can also be designed to replicate the behaviour of an iterative reconstruction by interlacing nonlinear convolutional filters with an operation that enforces closeness of the current image estimate to the acquired data^{15,16,18,19}, similar to the data fidelity step in an iterative shrinkage-thresholding algorithm²⁰. A recent study²¹ showed that neural network architectures that incorporate such an operation generalize better to different undersampling rates. In contrast, generic architectures that are solely based on convolutional layers can even lead to deteriorated image quality when the undersampling factor is increased, although one would expect the reconstruction result to improve when more information is available. The adversarial approach for training MR reconstruction networks^{22,23} is usually aimed at improving perceptual reconstruction quality, such as image sharpness. Typically, this is achieved at the expense of reconstruction normalized root-mean-square error (nRMSE)²⁴, which is critical for flow quantification.

In this work, an approach based on the idea of deep variational neural networks¹⁵ is implemented for rapid 4D flow reconstruction, which is referred to as FlowVN hereafter. For comparison, the 3D variational network (VN) architecture as presented in ref. ¹⁵ was adapted for 4D flow data by using 3D filter banks operating on subsets of $xyzt$ -dimensional data, yielding the model that we refer to as HamVN. The FlowVN network architecture replicates 10 steps of an iterative image reconstruction, while allowing for learnable spatiotemporal filter kernels, activation functions and regularization weights in each iteration. It is demonstrated that, based on training performed with retrospectively undersampled data of healthy subjects, FlowVN can accurately reconstruct pathological flow in a stenotic aorta in 21 s. Moreover, an imaging study with healthy subjects demonstrates good agreement of reconstructions from prospective undersampling with parallel imaging reference measurements.

4D flow MRI reconstruction with FlowVN

The FlowVN architecture improves HamVN in the following ways: (1) linear activations are used instead of radial basis functions, (2) the network is conditioned on the sampling rate, (3) exponential weighting

¹Institute for Biomedical Engineering, University and ETH Zurich, Zurich, Switzerland. ²These authors contributed equally: Valery Vishnevskiy, Jonas Walheim. ✉e-mail: vishnevskiy@biomed.ee.ethz.ch

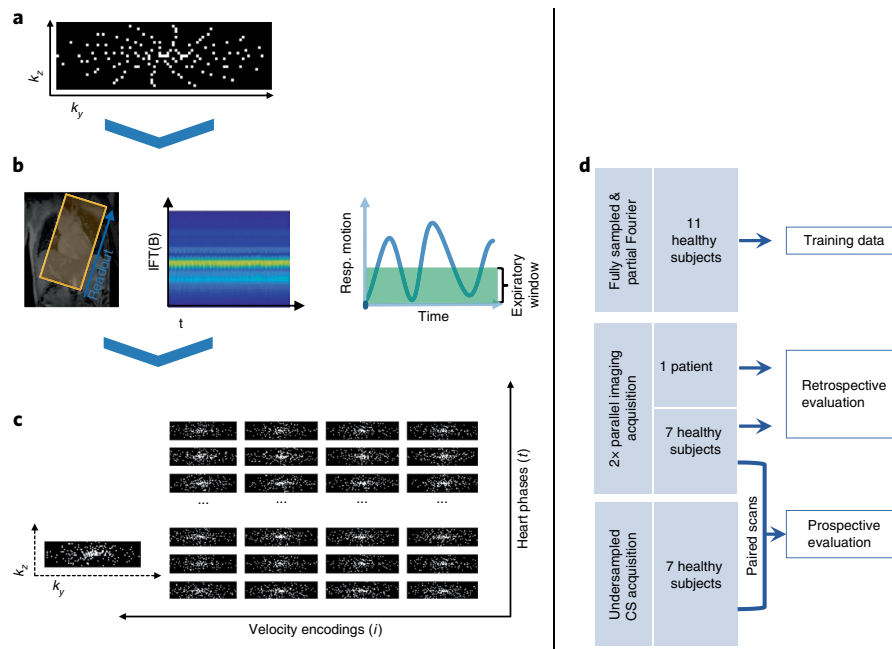


Fig. 1 | Breathing-resolved 4D flow data acquisition. **a**, Data are sampled using a Cartesian pseudo-radial tiny golden angle sampling pattern. **b**, Time-resolved data are inverse Fourier transformed (IFT) in the readout dimension, binned to end expiration using a combination of principal component analysis, low-pass filtering and coil clustering. **c**, Acquired data are sorted according to heart phase and velocity encoding. **d**, The datasets used during training and evaluation. To conduct prospective evaluation, seven healthy volunteers underwent accelerated compressed sensing (CS) and reference parallel imaging (PI) acquisitions during the same scan session.

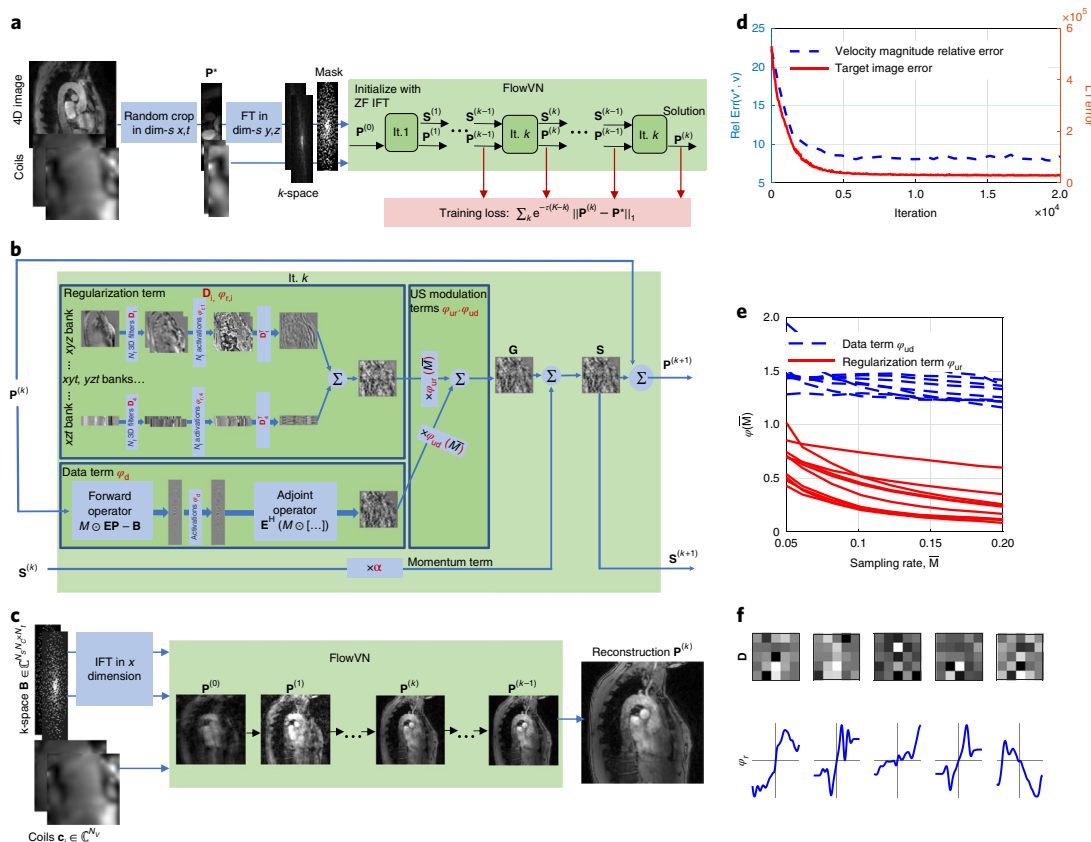


Fig. 2 | FlowVN architecture and training. **a**, Structure of FlowVN and its training strategy which uses reduced field-of-view data. **b**, Single unrolled iteration block consisting of convolutional filtering, data fidelity, undersampling (US) modulation and gradient descent momentum terms. Tunable parameters are highlighted in red. **c**, FlowVN at inference time yields 4D image reconstruction. **d**, Target training image error and velocity magnitude error in the aorta evaluated using training data. **e**, Data and gradient term weighting functions shown for each of 10 layers. **f**, Exemplary slices of 3D xyz filters and their corresponding activation functions at layer 5.

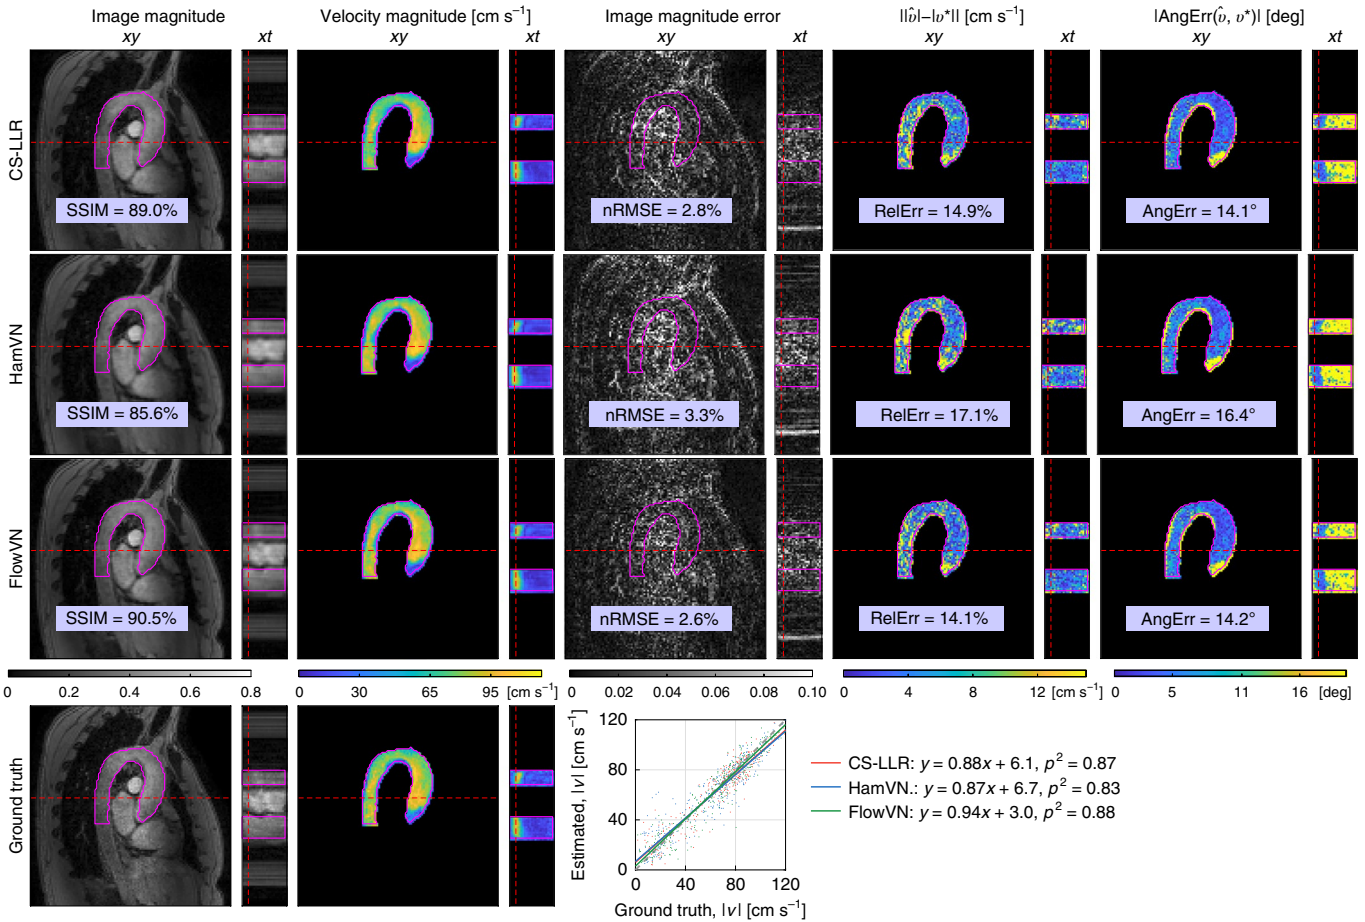


Fig. 3 | Reconstruction results on retrospectively undersampled data. Image magnitudes and estimated 4D velocity magnitude maps on retrospectively 14× undersampled data from a healthy volunteer. Corresponding slice locations are illustrated with red dashed lines that indicate cross-sections of the aorta and systolic peak. Bottom: scatter plot of velocity magnitude over the manually segmented aorta (contour shown in magenta throughout) together with correlation analysis ($y = ax + b$).

of intermediate layers is used as regularization, (4) real and imaginary parts of the signal are filtered by shared weights, (5) momentum is considered during gradient descent (GD) unrolling and (6) the data term allows tunable activation functions. The network is trained for a wide range of acceleration factors by allowing acceleration-dependent weighting of data consistency and filtering steps.

As illustrated in Fig. 1a, for each velocity encoding direction, the k -space data are acquired using a Cartesian golden angle sampling strategy, yielding variable density undersampling patterns in k -space. The signal of a total of 28 physical coils is compressed into 5 virtual coils via clustering²⁵. The samples are then sorted into respiratory bins and data in the end-expiratory bin are used for reconstruction.

A deep variational network can be seen as a differentiable sequence of an unrolled numerical optimization scheme. To enable learning, such a sequence is then relaxed by allowing tunable filter weights and activation functions. As described in the Methods, we unroll $K=10$ steps of a gradient descent with momentum governed by a scalar $\alpha^{(k)}$:

$$\mathbf{S}^{(k+1)} \leftarrow \alpha^{(k+1)}\mathbf{S}^{(k)} + \mathbf{G}^{(k)} \quad (1)$$

$$\mathbf{P}^{(k+1)} \leftarrow \mathbf{P}^{(k)} - \mathbf{S}^{(k+1)} \quad (2)$$

At each k th layer $\mathbf{P}^{(k)}$, the current complex-valued spatiotemporal image estimate is represented, while $\mathbf{S}^{(k)}$ maintains a running average of update steps. The update step $\mathbf{G}^{(k)}$ consists of the data

consistency and regularization terms (see Methods and Supplementary Algorithm 1 for details), which are weighted according to the sampling rate $\bar{M} = 1/R$ (R is the acceleration factor) via tunable activation functions for undersampling data $\phi_{\text{ud}}^{(k)}(\bar{M})$ and undersampling regularization $\phi_{\text{ur}}^{(k)}(\bar{M})$, respectively. The data consistency term modulates the k -space data residuals via an activation function and maps them back to the image space via a conjugate imaging operator. The regularization term at each layer contains 3D filters grouped into four banks, where each bank performs convolutions in three dedicated dimensions, namely xyz , xyt , xzt and yzt , therefore avoiding costly 4D convolutions. To avoid overfitting, we assume shared filters and activation functions that operate on real and imaginary components of the image. Note that both data and regularization terms do not assume correlations between real and imaginary parts of the signal, as highlighted in Fig. 2b.

The image estimate $\mathbf{P}^{(k)}(\mathbf{B}, \Theta)$ of the final layer can be then seen as a function of the k -space samples \mathbf{B} and network parameters Θ . To tune the network parameters Θ we minimize the layer-wise exponentially weighted ℓ_1 image reconstruction loss:

$$\min_{\Theta} \mathbb{E}_{\{\mathbf{B}, \mathbf{P}^*\} \sim \mathcal{T}} \sum_{k=1}^K e^{-\tau(K-k)} \|\mathbf{P}^{(k)}(\mathbf{B}, \Theta) - \mathbf{P}^*\|_1 \quad (3)$$

over the retrospectively undersampled training dataset \mathcal{T} , where \mathbf{P}^* is the ground-truth image. Layer weighting is controlled by

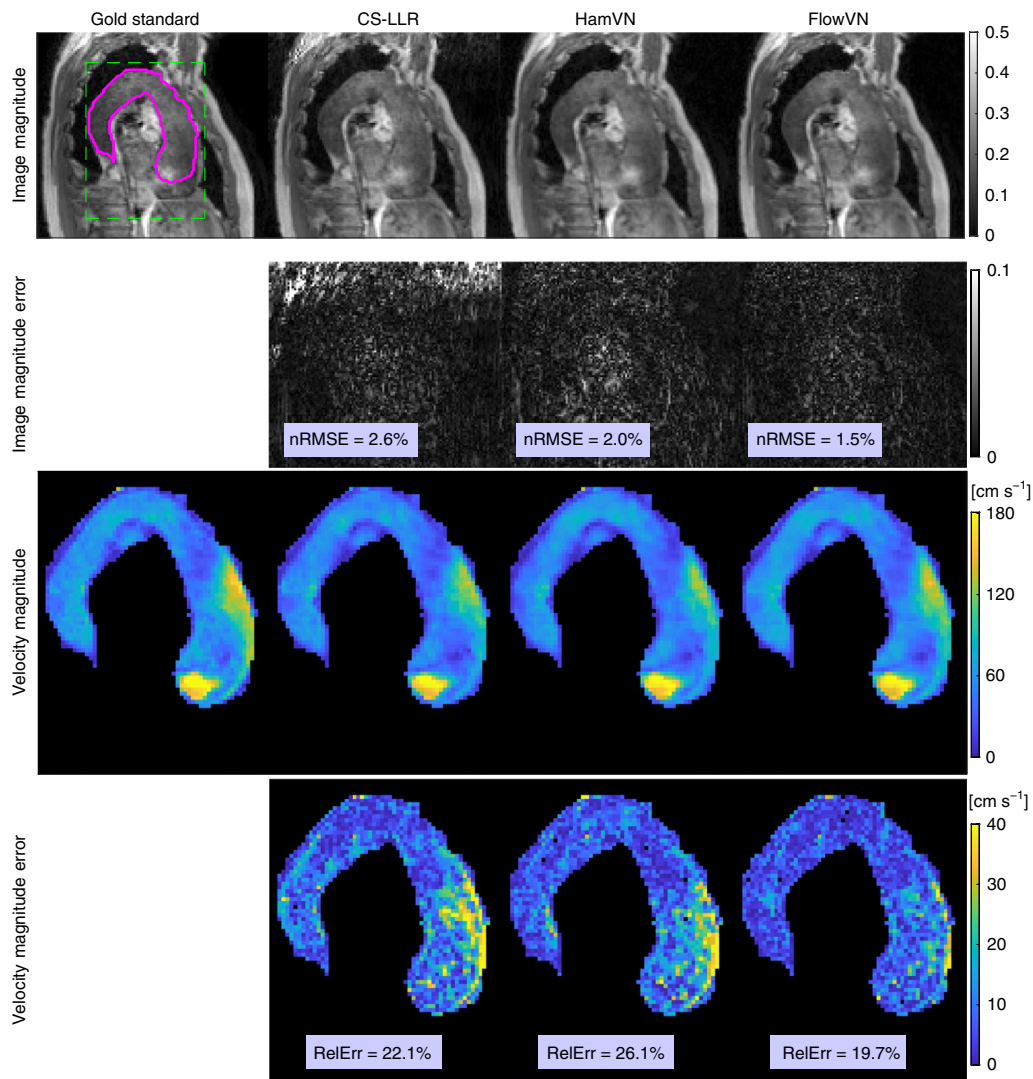


Fig. 4 | Retrospective reconstruction of the data from a patient with abnormal flow pattern. Reconstruction results of 10× retrospectively undersampled patient data shown at systolic peak flow. Manual aorta segmentation and the field of view are shown with magenta and green dashed lines, respectively.

parameter $\tau \geq 0$: when $\tau \approx 0$, the reconstruction error is penalized equally across layers, therefore gradients of network parameters have lower variance during stochastic optimization, yielding faster convergence. On the contrary, when $\tau \rightarrow +\infty$, only reconstruction at the final layer, $\mathbf{P}^{(K)}$, is minimized, which improves fitting accuracy on the training data. It is worth mentioning that τ controls the trade-off between training the reconstruction residual and network regularity. Similar to Landweber iterations^{15,26} and deep supervision²⁷, such implicit regularization penalizes irregular representations at intermediate layers and favours networks that can provide fast reconstruction. We propose to initialize τ with zero and then gradually increase it according to the training schedule (see Methods).

To demonstrate the validity of our approach, we note that the extracted velocity magnitude error in the aorta decreases simultaneously with the target reconstruction error during training, as shown in Fig. 2d, indicating that the target ℓ_1 image error is a valid training surrogate. It can be seen from Fig. 2e that the regularization term is suppressed for lower acceleration factors R (higher sampling rate \bar{M}). A subset of learned FlowVN parameters Θ is shown in Fig. 2f, illustrating that learned convolutions perform direction-dependent filtering.

Retrospective and prospective evaluation

Reconstructed image magnitudes (for a single velocity-encoding component), estimated velocity magnitudes and their errors

of healthy volunteer data for acceleration factor $R=14$ are shown in Fig. 3 for retrospectively undersampled data. Compared to CS-LLR and HamVN, the proposed FlowVN provides better reconstruction accuracy in terms of image magnitude and velocities. Scatter plot and correlation analysis further suggest that the velocity magnitude image that is estimated via FlowVN is in better agreement with the ground truth. As shown in Supplementary Table 1, these observations extend to other acceleration factors R (6–22), as tested on seven healthy volunteers.

Figure 4 indicates that FlowVN can accurately reconstruct the jet at the inlet section of the aorta for a patient with a pathological aortic valve.

The prospective undersampling acquisition results are reported in Fig. 5a,b: peak velocities and peak flow estimated using CS-LLR and FlowVN are in good agreement with PI reconstruction, while HamVN systematically underestimates the velocity magnitudes. Moreover, the correlation analysis shown in Fig. 5d reveals high correlation between CS-LLR and FlowVN velocity estimates. In contrast, HamVN shows systematic velocity underestimation compared to CS-LLR.

The exemplary reconstruction time for typical four-point velocity encoded images reported in Table 1 shows that the proposed FlowVN is 30 times faster than CS-LLR reconstruction.

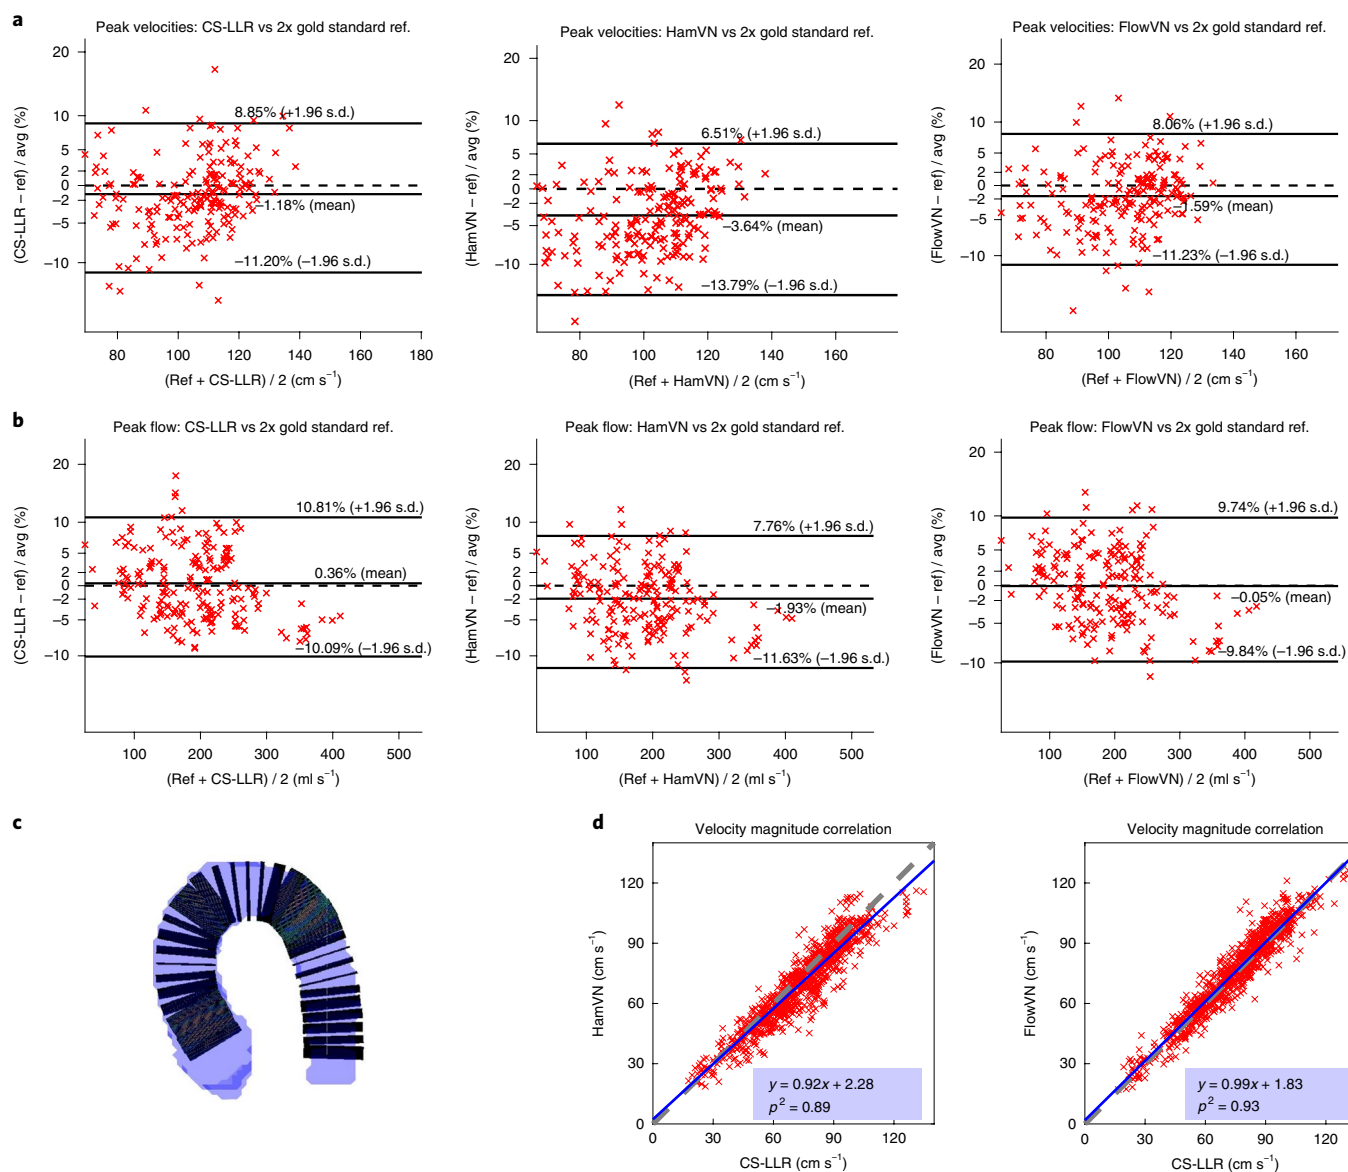


Fig. 5 | Quantitative flow evaluation of reconstruction methods on prospectively undersampled data ($12.4 \leq R \leq 13.8$) from seven healthy volunteers. **a–c**, Bland-Altman analyses of peak velocities (**a**) and peak flow (**b**) assessed over manually segmented aorta slices, as illustrated in **c**. Black dashed lines indicate zero difference between corresponding measurements and solid lines show mean and standard deviation values. **d**, Correlation analysis of aortic velocity magnitude estimates by learned variational network (VN) architectures and locally low rank (CS-LLR) reconstruction. The blue line shows the linear regression fit and the grey dashed line corresponds to $y = x$.

Discussion

Practical learning-based image reconstruction can be traced back to dictionary learning methods^{28,29}, where prior information is learned from image patches and then used as a sparsity-inducing regularizer for iterative reconstruction. Such an approach yields orders-of-magnitude longer reconstruction times compared to modern deep-learning methods. A straightforward application of deep artificial neural networks has been suggested for learning reconstruction as a regression from the k -space¹⁷ or zero-filled reconstructions³⁰ directly into the image space. Although tempting, such an approach might be unjustified, because the k -space and zero-filling artefacts have a global dependence on image intensities. The advent of effective automatic differentiation systems^{31,32} revitalized the idea of unrolling³³ and relaxing numerical schemes that can solve the original reconstruction problem. Following this approach, a number of deep neural network architectures were pro-

posed^{15,18,34,35} that disentangle image acquisition and image prior models. Unrolling gradient descent reconstruction with tunable filters and activation functions yields the HamVN architecture proposed by Hammernik¹⁵. One advantage of VNs is that, compared to other deep architectures, they have a relatively limited number of free parameters to tune, so they are less susceptible to overfitting.

In this work we have further developed the VN architecture^{15,36,37} to accommodate high-performance undersampled 4D flow reconstruction with limited training. Namely, we avoid exponential model complexity growth by avoiding 4D convolutions and by using separable 3D convolutions that are shared for real and imaginary parts of the image. Furthermore, in contrast to the original HamVN¹⁵, we train our FlowVN for a wide range of undersampling factors by allowing the regularization term to depend on them. As illustrated in Fig. 2e, regularization scaled by φ_{ur} decreases as more samples are available, while the data term φ_{ud} stays constant for most layers. Such

Table 1 | Model complexities and typical reconstruction time for 4D flow reconstruction

Method	Reconstruction time	No. of parameters
CS-LLR	10 min 24 s	2
HamVN	89 s	62,742
FlowVN	21 s	63,583

Typical reconstruction times are shown for four-point velocity encoded data compressed to five virtual coils and reconstructed on a $113 \times 113 \times 25$ grid. CS-LLR was executed on a six-core Intel CPU; FlowVN and HamVN were implemented in Tensorflow and evaluated on a NVIDIA Titan RTX system.

conditioning allows network training on a larger variety of artefacts and is necessary in practice, because, for a given fixed acquisition time, the precise value of the undersampling factor is not known a priori and depends on breathing and cardiac motion patterns. We hypothesize that the wide range of acceleration factors that were used simultaneously to train the FlowVN provided a diverse collection of aliasing artefacts and enabled robust learning on a remarkably limited training set of 11 subjects. The exponential weighting of the layer-wise reconstruction loss (equation (3)) further regularized FlowVN parameters by penalizing the nonlinear behaviour presented in HamVN reconstructions. Supplementary Fig. 1 and Supplementary Table 2 provide quantification of the reconstruction accuracy effects attributed to the modifications proposed with FlowVN. In particular, modifications to the network architecture result in a model that can better adapt to data and yield higher accuracy for retrospectively undersampled experiments, while the proposed exponential weighting of the training loss improves accuracy of the prospective evaluation, which indicates better generalization ability. It is worth noting that FlowVN has only 1% more tunable parameters than HamVN (Table 1) while improving the reconstruction nRMSE by 23% (averaged over acceleration factors, as given in Supplementary Fig. 1). We note that 4D flow MRI greatly benefits from using coil information during reconstruction (Supplementary Table 2). Accordingly, comparison with single-coil reconstruction networks¹⁸ has limited benefit.

The proposed FlowVN is a learning-based approach for reconstructing undersampled 4D flow MRI data in under a minute. For fixed reconstruction accuracy, FlowVN enables higher acceleration factors (12% improvement compared to the CS-LLR image nRMSE at $R=16$) and does not introduce significant bias of peak flow estimates. The proposed reconstruction is 30 times faster than state-of-the-art CS-LLR and 4.2 times faster than HamVN due to the use of linear activation functions rather than radial basis functions, which require computation of pairwise distances between control knots and image intensities. It is worth noting that FlowVN demonstrates high generalization ability, being able to preserve patient pathologies that were not present in the training data.

Methods

Compressed sensing 4D flow reconstruction. Phase contrast (PC) MRI encodes flow velocity $\mathbf{v}(\mathbf{r}, t) \in \mathbb{R}^3$ at spatial location \mathbf{r} during cardiac phase t ($1 \leq t \leq N_t$) according to the following equation:

$$\rho_i(\mathbf{r}, t) = \rho_0(\mathbf{r}, t) \exp\left(j\pi \frac{(\Phi \mathbf{v}(\mathbf{r}, t))_i}{v_{\text{enc}}}\right) \quad (4)$$

where v_{enc} is the velocity corresponding to a phase of $\pm\pi$, y_i ($i=0, \dots, 3$) are the encoded velocity vector components. The four-point velocity encoding matrix is given as

$$\Phi = \begin{bmatrix} 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (5)$$

Therefore, flow velocity \mathbf{v} can be calculated from the phase difference of reconstructed PC images ρ_i .

Let $\rho_{it} \in \mathbb{C}^{N_v}$ be a discretized image on an $N_x \times N_y \times N_z = N_v$ grid corresponding to a cardiac phase t and velocity encoding i . Assuming Cartesian sampling on a regular $N_1 \times N_2 \times N_3 = N_s$ grid, the Fourier transform $\mathbf{F} \in \mathbb{C}^{N_s \times N_s}$, and N_c coil sensitivity maps $\mathbf{W}_k \stackrel{\text{def}}{=} \text{diag}(\mathbf{c}_k) \in \mathbb{C}^{N_s \times N_s}$, define the spatial encoding operator $\mathbf{E} \in \mathbb{C}^{N_s N_c \times N_v}$:

$$\mathbf{E}\rho \stackrel{\text{def}}{=} \left[(\mathbf{F}\mathbf{W}_1\rho)^T, \dots, (\mathbf{F}\mathbf{W}_{N_c}\rho)^T \right]^T \in \mathbb{C}^{N_s N_c} \quad (6)$$

Considering a single velocity encoded image sequence, let $\mathbf{P} \in \mathbb{C}^{N_v \times N_t}$ and $\mathbf{B} \in \mathbb{C}^{N_s N_c \times N_t}$ be stacked column vectors of signals ρ and zero-filled k -space samples, respectively, while $\mathbf{M} \in \{0, 1\}^{N_s N_c \times N_t}$ defines the undersampling mask. Iterative image reconstruction methods seek for a maximum a posteriori (MAP) solution defined by the following optimization problem:

$$\hat{\mathbf{P}}_{\text{MAP}} = \underset{\mathbf{P}}{\text{argmin}} \frac{1}{2} \|\mathbf{M} \odot (\mathbf{E}\mathbf{P} - \mathbf{B})\|_F^2 + \mathcal{R}(\mathbf{P}) \quad (7)$$

where the regularization term \mathcal{R} enforces prior assumptions about image regularities. Here, we consider the LLR regularization¹⁴ to leverage image correlations among cardiac phases:

$$\mathcal{R}_{\text{LLR}}(\mathbf{P}) = \lambda_{\text{LLR}} \sum_{i \leq N_{\text{patch}}} \|\mathbf{T}_i \mathbf{P}\|_* \quad (8)$$

where $\mathbf{T}_i \in \{0, 1\}^{p \times p \times p}$ is the corresponding $p \times p \times p$ patch extraction operator, yielding N_{patch} patches, and $\|\cdot\|_*$ is the nuclear norm. For LLR regularization, the optimization problem (7) is convex and can be efficiently solved using operator splitting techniques such as the fast iterative shrinkage-thresholding algorithm (FISTA)²⁰.

FlowVN training. We employ a $K=10$ layer VN and perform 5×10^4 iterations of the ADAM algorithm (learning rate 10^{-3} , $\beta_1=0.85$, $\beta_2=0.98$, batch size of three) for training, during which we continually adjust $\tau = i_{\text{opt}} \times 10^{-3}$, where i_{opt} is the iteration number. On every layer, each 3D filter bank contains $N_f=8$ filters of size $n_x=5$ voxels. Activation functions $\phi\{\cdot\}$ are parametrized by $N_{\text{knots}}=91$ control knots with spacing $\omega=0.17$:

$$\phi\{h\} = (1 - h/\omega + [h/\omega])\phi_{[h/\omega]} + (h/\omega - [h/\omega])\phi_{[h/\omega]+1} \quad (9)$$

with gradients provided by the following formulae:

$$\begin{aligned} \frac{\partial \phi}{\partial h} \{h\} &= 1_{i \leq h \leq i+1} (1 - h + [h]) + 1_{i-1 \leq h \leq i} (h - [h]) \\ \frac{\partial \phi}{\partial h} \{h\} &= \phi_{[h]+1} - \phi_{[h]} \end{aligned} \quad (10)$$

The acquired zero-filled k -space \mathbf{B} with undersampling mask \mathbf{M} was normalized by $\frac{\|\mathbf{M}\|_1}{\|\mathbf{B}\|_1}$.

To enable backpropagation to be carried out with limited GPU memory, we employ spatiotemporal equivariance of the convolution and exploit the fact that the k -space is fully sampled in the readout dimension k_x for Cartesian acquisitions. Therefore, to draw a training sample, we perform random cropping of width w_x and w_t in dimensions x and t , respectively and simulate Fourier encoding in the $k_y k_z$ dimensions, as illustrated in Fig. 2a. The network was implemented using the Tensorflow framework³². Fully sampled and partial Fourier acquisition data from 11 healthy volunteers were used during training.

In vivo data acquisition. As illustrated in Fig. 1, we used eleven subjects for network training and seven healthy subjects and one patient for evaluation. All in vivo work was performed on receiving written informed consent from the subjects and according to local ethics regulations.

The training datasets comprised 4D flow data measured in the aorta of eleven healthy subjects, nine of them fully sampled and two acquired with partial Fourier approach³⁸ (factor 0.75×0.75).

For evaluation, data in the ascending aorta of seven healthy subjects were acquired on a 3T Philips Ingenia system (Philips Healthcare, Best, the Netherlands) using a Cartesian four-point referenced phase-contrast gradient-echo sequence with an encoding velocity of $v_{\text{enc}} = 150 \text{ cm s}^{-1}$, spatial resolution of $2.5 \times 2.5 \times 2.5 \text{ mm}^3$, $T_E = 3.3 \text{ ms}$, $T_R = 4.9 \text{ ms}$, 25 cardiac phases and flip angle of 8° . Examinations for each of the seven healthy subjects comprised a standard navigator-gated twofold accelerated parallel imaging⁵ examination for reference and a CS acquisition with an acceleration factor of $R = 12.4\text{--}13.8$, using a Cartesian pseudo-radial golden angle sampling pattern³⁹ and data-driven respiratory motion detection, as in ref. ¹¹. Only data in expiration were kept for reconstruction, as shown in Fig. 1.

To evaluate the reconstruction accuracy for pathological anatomy, 4D flow data were acquired in a single patient with dilation of the ascending aorta, combined aortic stenosis and regurgitation due to a bicuspid aortic valve on a 3T Philips

Ingenia system (Philips Healthcare), using a navigator-gated twofold-accelerated parallel imaging⁵ scan.

A receiver coil with 28 channels was used for acquisition; these channels were reduced to five channels using coil compression⁴⁰. Coil sensitivity maps were estimated with ESPiRiT⁴¹. Concomitant field correction was applied to the signal phase according to ref.⁴² and eddy currents were corrected with a third-order polynomial model fitted to stationary tissue^{43,44}.

Evaluation. We compared the proposed FlowVN to the state-of-the-art compressed sensing LLR-regularized (8) reconstruction¹⁴ and the variational network from Hammernik and colleagues¹⁵, which we refer to as HamVN. The LLR implementation from the Berkeley Advanced Reconstruction Toolbox (BART)⁴⁵ was used with a patch size of $p=8$ and a maximum number of optimization iterations of 80. The optimal value of the regularization parameter, $\lambda_{\text{LLR}}=2.06$, was chosen via a grid search to minimize the reconstructed flow field residual $\|\tilde{v} - v^*\|_2$ averaged over the manually segmented aorta on the retrospectively 12× undersampled acquisition. Because the original VN¹⁵ was proposed for magnitude reconstruction of 2D and 3D data, we introduced the following modifications for the presented 4D flow evaluation: (1) 3D filters were grouped into four banks as in FlowVN (Supplementary algorithm (1)), (2) the ℓ_1 -norm of reconstruction was optimized (that is, equation (3) with $\tau \rightarrow +\infty$). We refer to this architecture as ‘HamVN’; the number of network layers, filters and control knots were the same as in FlowVN.

Retrospective study. For simulated retrospective undersampling experiments, we used 2×PI data and simulated a pseudo-radial golden angle sampling pattern³⁹ with acceleration factors of 6–22.

For each undersampling factor we evaluated the nRMSE of the image magnitude, the relative error (RelErr) of velocity magnitudes inside the aorta and the angular error (AngErr) of the estimated velocity vectors:

$$\begin{aligned} \text{nRMSE}(\mathbf{a}, \mathbf{a}^*) &= \sqrt{\frac{\sum_i^N (a_i - a_i^*)^2}{N \max_j (a_j^*)^2}} \\ \text{RelErr}(\mathbf{a}, \mathbf{a}^*) &= \frac{\|\mathbf{a} - \mathbf{a}^*\|_2}{\|\mathbf{a}^*\|_2} \\ \text{AngErr}(\mathbf{u}, \mathbf{v}) &= \arccos\left(\frac{\langle \mathbf{u}, \mathbf{v} \rangle}{\|\mathbf{u}\|_2 \|\mathbf{v}\|_2}\right) \end{aligned} \quad (11)$$

Additionally, we report the structural similarity index (SSIM)⁴⁶ with $\sigma_{\text{SSIM}}=1.5$ on the reconstructed magnitude images.

Prospective study. Using manual aorta segmentations we computed flow over cross-sections of the aorta by integrating velocity components projected onto the cross-section normal. The peak flow was then defined as the maximal flow over cardiac phases for a given cross-section. Moreover, we calculated the peak through-plane velocity, defined as the maximum velocity projection across cross-sections of the aorta over cardiac phases.

To quantify agreement with the reference 2×PI reconstruction, we performed Bland–Altman analysis⁴⁷ of peak flow and peak through-plane velocities.

Reporting Summary. Further information on research design is available in the Nature Research Reporting Summary linked to this article.

Data availability

The code for the network training and inference used in this study as well as network weights are available online from CodeOcean together with volunteer data: <https://codeocean.com/capsule/0115983/tree>⁴⁸. The code for analysis is available on CodeOcean from <https://codeocean.com/capsule/2587940/tree>⁴⁹.

Received: 4 October 2019; Accepted: 11 March 2020;

Published online: 13 April 2020

References

- Markl, M., Frydrychowicz, A., Kozerke, S., Hope, M. & Wieben, O. 4D flow MRI. *J. Magn. Res. Imag.* **36**, 1015–1036 (2012).
- Feinberg, D., Hale, J., Watts, J., Kaufman, L. & Mark, A. Halving MR imaging time by conjugation: demonstration at 3.5kG. *Radiology* **161**, 527–531 (1986).
- Szarf, G. et al. Zero filled partial Fourier phase contrast MR imaging: in vitro and in vivo assessment. *J. Magn. Reson. Imag.* **23**, 42–49 (2006).
- Walheim, J., Gotschy, A. & Kozerke, S. On the limitations of partial Fourier acquisition in phase-contrast MRI of turbulent kinetic energy. *Magn. Reson. Med.* **81**, 514–523 (2019).
- Pruessmann, K., Weiger, M., Scheidegger, M. & Boesiger, P. SENSE: sensitivity encoding for fast MRI. *Magn. Reson. Med.* **42**, 952–962 (1999).
- Wiesinger, F., Boesiger, P. & Pruessmann, K. P. Electrodynamics and ultimate SNR in parallel MR imaging. *Magn. Reson. Med.* **52**, 376–390 (2004).
- Lustig, M., Donoho, D. & Pauly, J. M. Sparse MRI: the application of compressed sensing for rapid MR imaging. *Magn. Reson. Med.* **58**, 1182–1195 (2007).
- Kim, D. et al. Accelerated phase-contrast cine MRI using k - t SPARSE-SENSE. *Magn. Reson. Med.* **67**, 1054–1064 (2012).
- Valvano, G. et al. Accelerating 4D flow MRI by exploiting low-rank matrix structure and hadamard sparsity. *Magn. Reson. Med.* **78**, 1330–1341 (2017).
- Bollache, E. et al. k - t accelerated aortic 4D flow MRI in under two minutes: feasibility and impact of resolution, k -space sampling patterns, and respiratory navigator gating on hemodynamic measurements. *Magn. Reson. Med.* **79**, 195–207 (2018).
- Walheim, J., Dillinger, H. & Kozerke, S. Multipoint 5D flow cardiovascular magnetic resonance—accelerated cardiac- and respiratory-motion resolved mapping of mean and turbulent velocities. *J. Cardiovasc. Magn. Res.* **21**, 42 (2019).
- Ma, L. E. et al. Aortic 4D flow MRI in 2 minutes using compressed sensing, respiratory controlled adaptive k -space reordering, and inline reconstruction. *Magn. Reson. Med.* **81**, 3675–3690 (2019).
- Rich, A. et al. A Bayesian approach for 4D flow imaging of aortic valve in a single breath-hold. *Magn. Reson. Med.* **81**, 811–824 (2019).
- Zhang, T., Pauly, J. M. & Levesque, I. R. Accelerating parameter mapping with a locally low rank constraint. *Magn. Reson. Med.* **73**, 655–661 (2015).
- Hammernik, K. et al. Learning a variational network for reconstruction of accelerated MRI data. *Magn. Reson. Med.* **79**, 3055–3071 (2018).
- Mardani, M. et al. Deep generative adversarial neural networks for compressive sensing MRI. *IEEE Trans. Med. Imag.* **38**, 167–179 (2018).
- Zhu, B., Liu, J. Z., Cauley, S. F., Rosen, B. R. & Rosen, M. S. Image reconstruction by domain-transform manifold learning. *Nature* **555**, 487–492 (2018).
- Schlemper, J., Caballero, J., Hajnal, J. V., Price, A. N. & Rueckert, D. A deep cascade of convolutional neural networks for dynamic MR image reconstruction. *IEEE Trans. Med. Imag.* **37**, 491–503 (2017).
- Maier, A. K. et al. Learning with known operators reduces maximum error bounds. *Nat. Mach. Intell.* **1**, 373–380 (2019).
- Beck, A. & Teboulle, M. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imag. Sci.* **2**, 183–202 (2009).
- Antun, V., Renna, F., Poon, C., Adcock, B. & Hansen, A. C. On instabilities of deep learning in image reconstruction—does AI come at a cost? Preprint at <https://arxiv.org/pdf/1902.05300.pdf> (2019).
- Yang, G. et al. DAGAN: deep de-aliasing generative adversarial networks for fast compressed sensing MRI reconstruction. *IEEE Trans. Med. Imag.* **37**, 1310–1321 (2017).
- Quan, T., Nguyen-Duc, T. & Jeong, W. Compressed sensing MRI reconstruction using a generative adversarial network with a cyclic loss. *IEEE Trans. Med. Imag.* **37**, 1488–1497 (2018).
- Narnhofer, D., Hammernik, K., Knoll, F. & Pock, T. Inverse GANs for accelerated MRI reconstruction. *Proc. SPIE* **11138**, 111381A (2019).
- Zhang, S., Block, K. & Frahm, J. Magnetic resonance imaging in real time: advances using radial FLASH. *J. Mag. Res. Imag.* **31**, 101–109 (2010).
- Landweber, L. An iteration formula for Fredholm integral equations of the first kind. *Am. J. Math.* **73**, 615–624 (1951).
- Liu, Y. & Lew, M. S. Learning relaxed deep supervision for better edge detection. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition* 231–240 (CVPR, 2016).
- Ravishanker, S. & Bresler, Y. MR image reconstruction from highly undersampled k -space data by dictionary learning. *IEEE Trans. Med. Imag.* **30**, 1028–1041 (2010).
- Caballero, J., Price, A. N., Rueckert, D. & Hajnal, J. V. Dictionary learning and time sparsity for dynamic MR data reconstruction. *IEEE Trans. Med. Imag.* **33**, 979–994 (2014).
- Lee, D., Yoo, J. & Ye, J. C. Deep residual learning for compressed sensing MRI. In *Proc. of IEEE Inter. Symp. on Bio. Imag.* 15–18 (2017).
- LeCun, Y. A theoretical framework for back-propagation. In *Proc. Connectionist Models Summer School* Vol. 1, 21–28 (Morgan Kaufmann, 1988).
- Abadi, M. et al. Tensorflow: a system for large-scale machine learning. In *Proceedings of the 12th USENIX conference on Operating Systems Design and Implementation* 265–283 (2016).
- Domke, J. Generic methods for optimization-based modeling. In *Proc. Int. Conf. on Artificial Intelligence and Statistics* Vol. 22, 318–326 (ACM, 2012).
- Sun, J. et al. Deep ADMM-Net for compressive sensing MRI. In *Proc. Int. Conf. on Neural Information Processing Systems* 10–18 (Curran Associates, 2016).
- Jin, K. H., McCann, M. T., Froustey, E. & Unser, M. Deep convolutional neural network for inverse problems in imaging. *IEEE Trans. Imag. Proc.* **26**, 4509–4522 (2017).
- Vishnevskiy, V., Sanabria, S. J. & Goksel, O. Image reconstruction via variational network for real-time hand-held sound-speed imaging. In *First International Workshop for Machine Learning for Medical Image Reconstruction* 120–128 (Springer, 2018).
- Vishnevskiy, V., Rau, R. & Goksel, O. Deep variational networks with exponential weighting for learning computed tomography. In *22nd International Conference on Medical Image Computing and Computer-Assisted Intervention Part VI*, Vol. 11769, 310–318 (LNCS, Springer, 2019).

38. Cuppen, J. & van Est, A. Reducing MR imaging time by one-sided reconstruction. *Magn. Reson. Imag.* **5**, 526–527 (1987).
39. Winkelmann, S., Schaeffter, T., Koehler, T., Eggers, H. & Doessel, O. An optimal radial profile order based on the golden ratio for time-resolved MRI. *IEEE Trans. Med. Imag.* **26**, 68–76 (2006).
40. Zhang, T., Pauly, J. M., Vasanawala, S. S. & Lustig, M. Coil compression for accelerated imaging with cartesian sampling. *Magn. Reson. Med.* **69**, 571–582 (2013).
41. Uecker, M. et al. ESPIRiT—an eigenvalue approach to autocalibrating parallel MRI: where SENSE meets GRAPPA. *Magn. Reson. Med.* **71**, 990–1001 (2014).
42. Bernstein, M. A. et al. Concomitant gradient terms in phase contrast MR: analysis and correction. *Magn. Reson. Med.* **39**, 300–308 (1998).
43. Busch, J., Giese, D. & Kozerke, S. Image-based background phase error correction in 4D flow MRI revisited. *J. MRI* **46**, 1516–1525 (2017).
44. Walker, P. G. et al. Semiautomated method for noise reduction and background phase error correction in MR phase velocity data. *J. Magn. Reson. Imag.* **3**, 521–530 (1993).
45. Tamir, J. I., Ong, F., Cheng, J. Y., Uecker, M. & Lustig, M. Generalized magnetic resonance image reconstruction using the Berkeley advanced reconstruction toolbox. In *ISMRM Workshop on Data Sampling and Image Reconstruction* (ISMRM, 2016).
46. Wang, Z., Bovik, A. C., Sheikh, H. R. & Simoncelli, E. P. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Imag. Proc.* **13**, 600–612 (2004).
47. Altman, D. G. & Bland, J. M. Measurement in medicine: the analysis of method comparison studies. *J. R. Stat. Soc.* **32**, 307–317 (1983).
48. Vishnevskiy, V., Walheim, J. & Kozerke, S. FlowVN: deep variational network for rapid 4D flow MRI reconstruction. *CodeOcean* <https://doi.org/10.24433/CO.0115983.v1> (2020).
49. Vishnevskiy, V., Walheim, J. & Kozerke, S. FlowVN: analysis. *CodeOcean* <https://doi.org/10.24433/CO.5994453.v1> (2020).

Acknowledgements

The authors acknowledge funding from the European Unions Horizon 2020 research and innovation programme under grant agreement no. 668039 and under EuroStars UNIFORM as well as funding of the Platform for Advanced Scientific Computing of the Council of the Federal Institutes of Technology (ETH Board), Switzerland.

Author contributions

J.W., V.V. and S.K. conceived the study. V.V. implemented the machine-learning reconstruction algorithms. J.W. conducted MR acquisition experiments and data preprocessing. J.W. and V.V. analysed experimental data under the supervision of S.K. All authors discussed the results and contributed to writing the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s42256-020-0165-6>.

Correspondence and requests for materials should be addressed to V.V.

Reprints and permissions information is available at www.nature.com/reprints.

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

© The Author(s), under exclusive licence to Springer Nature Limited 2020

Reporting Summary

Nature Research wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Research policies, see [Authors & Referees](#) and the [Editorial Policy Checklist](#).

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a Confirmed

- The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement
- A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
- The statistical test(s) used AND whether they are one- or two-sided
Only common tests should be described solely by name; describe more complex techniques in the Methods section.
- A description of all covariates tested
- A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
- A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
- For null hypothesis testing, the test statistic (e.g. F , t , r) with confidence intervals, effect sizes, degrees of freedom and P value noted
Give P values as exact values whenever suitable.
- For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
- For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
- Estimates of effect sizes (e.g. Cohen's d , Pearson's r), indicating how they were calculated

Our web collection on [statistics for biologists](#) contains articles on many of the points above.

Software and code

Policy information about [availability of computer code](#)

Data collection

In-vivo measurements were performed on Philips 3.0T MR system; two-fold undersampled reference measurements using standard 4D Flow MRI protocol with parallel imaging incl. off-line image reconstruction using MRRecon (GyroTools LLC, Zurich, www.gyrotools.com).

Data analysis

MATLAB 2017a (Mathworks, Natick, MA, USA), Berkeley advanced reconstruction toolbox (BART), Python 3.5.2 with NumPY 1.15.1 and Tensorflow 1.14 were used to implement algorithms and data analysis

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors/reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Research [guidelines for submitting code & software](#) for further information.

Data

Policy information about [availability of data](#)

All manuscripts must include a [data availability statement](#). This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A list of figures that have associated raw data
- A description of any restrictions on data availability

Data will be publicly available before publication

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

- Life sciences Behavioural & social sciences Ecological, evolutionary & environmental sciences

Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size	18 healthy subjects and 1 patient who volunteered to participate in MRI exams upon written informed consent according to institutional and ethics guidelines
Data exclusions	No data were excluded during analysis
Replication	We split the retrospectively undersampled data into training and test sets that contained data from volunteers not presented in training and one patient with irregular flow in the aorta. We also included a prospective study, where undersampled signal was acquired on the scanner in volunteers that were not used during training.
Randomization	Healthy subjects were scanned using two-fold undersampled 4D Flow MRI protocol with parallel imaging (PI) to provide training data. Subjects that underwent both 2xPI and ~14x accelerated pseudo Golden angle acquisitions were used for both retrospective and prospective studies.
Blinding	Since quantitative accuracy of the reconstructed flow measurements are reported and hence blinding was not required.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involvement
<input checked="" type="checkbox"/>	<input type="checkbox"/> Antibodies
<input checked="" type="checkbox"/>	<input type="checkbox"/> Eukaryotic cell lines
<input checked="" type="checkbox"/>	<input type="checkbox"/> Palaeontology
<input checked="" type="checkbox"/>	<input type="checkbox"/> Animals and other organisms
<input type="checkbox"/>	<input checked="" type="checkbox"/> Human research participants
<input checked="" type="checkbox"/>	<input type="checkbox"/> Clinical data

Methods

n/a	Involvement
<input checked="" type="checkbox"/>	<input type="checkbox"/> ChIP-seq
<input checked="" type="checkbox"/>	<input type="checkbox"/> Flow cytometry
<input checked="" type="checkbox"/>	<input type="checkbox"/> MRI-based neuroimaging

Human research participants

Policy information about [studies involving human research participants](#)

Population characteristics	For training and analysis we used data from 18 healthy volunteers (age 30+-11) and 1 patient with dilation of the ascending aorta and combined aortic stenosis and regurgitation due to a bicuspid aortic valve.
Recruitment	Volunteers were recruited locally; the patient was recruited as part of routine examination at the University Hospital Zurich.
Ethics oversight	ETH Zurich

Note that full information on the approval of the study protocol must also be provided in the manuscript.