



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Unsupervised and Segmentation-Free Phase Unwrapping for 4D Flow MRI Using Deep Image Prior

Yuyang Ren¹, Zijian Zhou¹, and Peng Hu^{1,2} *

¹ School of Biomedical Engineering & State Key Laboratory of Advanced Medical Materials and Devices, ShanghaiTech University, Shanghai 201210, China

² Shanghai Clinical Research and Trial Center, Shanghai 201210, China
renyy2022@shanghaitech.edu.cn, zhouzj@shanghaitech.edu.cn,
hupeng@shanghaitech.edu.cn

Abstract. 4D flow MRI provides rich hemodynamic information by encoding blood-flow velocities in the image phase; however, image quality highly depends on the velocity-encoding (venc) setting—high venc reduces the velocity-to-noise ratio (VNR), whereas low venc improves VNR but increases the risk of phase wrapping, potentially introducing substantial bias in downstream flow quantification. We propose PUDIP-Flow, an unsupervised and segmentation-free phase-unwrapping method for aortic 4D flow MRI based on the deep image prior (DIP). The method formulates phase unwrapping as an optimization problem that enforces consistency with wrapped spatiotemporal phase gradients, and solves it by fitting a 3D U-Net to a single scan without paired training data. To reduce reliance on manual segmentation of regions of interest, we further introduce an automated weighting mask derived from phase-contrast MR angiography features and temporal variance, thereby suppressing contributions from stationary background tissue. Experiments on retrospective and prospective aortic datasets show that PUDIP-Flow improves unwrapping accuracy and robustness compared with established baselines, with particularly strong performance when manual segmentations are unavailable.

Keywords: 4D flow MRI · phase unwrapping · deep image prior · unsupervised learning

1 Introduction

4D flow magnetic resonance imaging (MRI) [10] provides time-resolved 3D blood-flow velocity fields and enables clinically relevant measurements such as net flow, flow rate, wall shear stress, and pressure differences [7]. Velocities are derived from the phase of acquired images, where the velocity range is determined by a prescribed velocity encoding (venc) corresponding to a phase interval of $(-\pi, \pi]$.

* Corresponding author: hupeng@shanghaitech.edu.cn

Lower venc typically yields higher velocity-to-noise ratio, but velocities exceeding $\pm \text{venc}$ cause phase wrapping and bias their estimates.

Beyond selecting an appropriate venc, prior work has proposed various unwrapping strategies, including dual-venc acquisitions [2, 11] and single-venc unwrapping methods leveraging smoothness priors and optimization formulations [3, 9, 16, 19]. Dual-venc can be effective but increases scan time, while performance of single-venc methods may degrade under severe wrapping or low SNR. Deep learning has also been explored for 4D flow phase unwrapping [4], but supervised training requires paired datasets that are difficult to obtain.

We propose PUDIP-Flow, an unsupervised phase-unwrapping method for aortic 4D flow MRI based on the deep image prior (DIP) [15, 18]. DIP leverages the implicit bias of convolutional networks as an image prior and can be fit to a single scan without external training data. Our contributions are threefold: (i) we cast 4D flow phase unwrapping as a gradient-domain optimization problem and solve it by fitting a 3D U-Net to a single scan, avoiding paired training data; (ii) we propose an automatic weighting mask derived from phase-contrast MR angiography (PCMRA) and temporal variance, removing the need for manual segmentation; and (iii) we validate PUDIP-Flow on retrospective and prospective aortic datasets spanning a wide range of venc and SNR, achieving improved accuracy and robustness over established baselines.

2 Methods

2.1 Problem formulation

In 4D flow MRI, each velocity-encoding direction $i \in \{x, y, z\}$ yields a flow-encoded (FE) image $\mathbf{I}_{\text{FE},i} = \mathbf{I}_{\text{FC}} \odot e^{j\Psi_i}$, where $\mathbf{I}_{\text{FC}} \in \mathbb{C}^N$ is the flow-compensated (FC) image, $\Psi_i \in \mathbb{R}^N$ is the true velocity-induced phase, $\odot$ denotes the element-wise product, and $N = N_t N_x N_y N_z$ spans the spatiotemporal volume. The acquired (wrapped) phase is

$$\Phi_i = \angle(\mathbf{I}_{\text{FE},i} \odot \mathbf{I}_{\text{FC}}^*) = \mathcal{W}(\Psi_i) \in (-\pi, \pi], \quad (1)$$

where $(\cdot)^*$ is complex conjugate, $\angle(\cdot)$ returns the principal-value phase, and $\mathcal{W}(x) = (x + \pi) \bmod 2\pi - \pi$ is the wrapping operator. The velocity along direction i is obtained by linearly scaling the phase:

$$\mathbf{v}_i = \frac{\text{venc}}{\pi} \Phi_i, \quad (2)$$

where venc (velocity encoding) sets the measurable range of $[-\text{venc}, +\text{venc}]$. Velocities exceeding this range cause the phase to wrap, producing aliasing artifacts.

Phase unwrapping seeks an integer field $\mathbf{k}_i \in \mathbb{Z}^N$ such that the true phase is recovered as

$$\Psi_i = \Phi_i + 2\pi\mathbf{k}_i. \quad (3)$$




Fig. 1. Pipeline of PUDIP-Flow. A 3D U-Net is optimized per scan using the proposed loss.

Most methods exploit Itoh’s condition [8]: when the true phase varies smoothly ($|\nabla \Psi_i| \leq \pi$), the wrapped gradient equals the true gradient,

$$\mathcal{W}(\nabla \Phi_i) = \nabla \Psi_i, \quad (4)$$

where ∇ is the discrete gradient operator. Eq. (4) underpins weighted least-squares [19], Poisson-equation [14], and graph-cut [5] solvers. Recently, deep image prior (DIP) [15] has shown strong performance for phase unwrapping in quantitative phase imaging [18], and we extend this idea to 4D flow MRI.

2.2 Phase unwrapping via DIP for 4D flow MRI

To recover the unwrapped phase Ψ_i for each velocity-encoding direction i using DIP, a randomly initialized tensor $\mathbf{z}_i$ is fed into a neural network f_{θ_i} parameterized by θ_i . Denoting the network output as the estimated unwrapped phase, we write

$$\hat{\Psi}_i = f_{\theta_i}(\mathbf{z}_i), \quad (5)$$

where $\mathbf{z}_i \in \mathbb{R}^{N_c N_x N_y N_z}$ and N_c is the number of initial feature channels of the network. Specifically, f_{θ_i} is a 3D U-Net that operates on the three spatial dimensions; the N_t temporal frames are treated as output channels, so that spatial regularity is enforced by the convolutional architecture while all time frames are estimated jointly.

The optimal network parameters are obtained by solving the following optimization problem:

$$\hat{\theta}_i = \arg \min_{\theta_i} \|\mathbf{M} \odot (\mathcal{W}(\nabla \Phi_i) - \nabla f_{\theta_i}(\mathbf{z}_i))\|_1, \quad (6)$$

where $\nabla : \mathbb{R}^N \rightarrow \mathbb{R}^{4N}$ denotes the 4D discrete spatiotemporal finite-difference operator and $\mathbf{M} \in \mathbb{R}^N$ is a voxel-wise weighting map that suppresses the influence of the stationary background. Motivated by Itoh’s continuity condition Eq. (4), we formulate the objective in the gradient domain: by forcing the gradient of the network output, $\nabla f_{\theta_i}(\mathbf{z}_i)$, to progressively match the wrapped measure-phase gradient $\mathcal{W}(\nabla \Phi_i)$, the network output $\hat{\Psi}_i = f_{\theta_i}(\mathbf{z}_i)$ is driven toward the true unwrapped phase Ψ_i .

In principle, $\mathbf{M}$ can be a segmentation mask that precisely defines the region of interest (ROI) for phase unwrapping. However, obtaining such a mask typically requires manual delineation and is time-consuming. Therefore, we propose an automated method for computing the weighting mask based on PCMRA. The PCMRA image is defined as

$$\mathbf{P} = |\mathbf{I}_{\text{FC}}| \odot \sqrt{\Phi_x^2 + \Phi_y^2 + \Phi_z^2}, \quad (7)$$

where $|\cdot| : \mathbb{C}^N \rightarrow \mathbb{R}^N$ computes the element-wise magnitude, and the weighting mask is then obtained as

$$\mathbf{M} = \mathbf{P} \odot \text{std}_t(\mathbf{P}), \quad (8)$$

where $\text{std}_t(\cdot)$ denotes the standard deviation along the temporal dimension. Stationary regions are thereby assigned lower weights because they exhibit both lower velocities and smaller temporal velocity variations, whereas regions with high velocities and large temporal variations receive higher weights. The complete pipeline of PUDIP-Flow is illustrated in Fig. 1.

Once the estimated phase map $\hat{\Psi}_i = f_{\hat{\theta}_i}(\mathbf{z}_i)$ is obtained, an additional post-processing step [12] is applied to ensure that the final phase map differs from the acquired phase by an integer multiple of 2π (Eq. (3)):

$$\hat{\Psi}'_i = f_{\hat{\theta}_i}(\mathbf{z}_i) + \mathcal{W}(\Phi_i - f_{\hat{\theta}_i}(\mathbf{z}_i)). \quad (9)$$

The final velocity in direction i is then computed as

$$\hat{\mathbf{v}}_i = \frac{\text{venc}}{\pi} \hat{\Psi}'_i. \quad (10)$$

2.3 Experimental setup

We evaluated PUDIP-Flow on aortic 4D flow MRI using retrospective and prospective experiments. The network used an initial feature width of $N_c = 128$; the random input was drawn i.i.d. from $\mathcal{N}(0, 1)$ and scaled by 0.1 and optimized for $K = 1000$ iterations.

A total of 27 healthy subjects were included. Seven cases were from a public database [17] (venc 150 cm/s, spatial resolution 2.5 mm isotropic, 25 temporal frames). Twenty cases were acquired using a self-gated dual-venc 4D flow sequence on a 3T scanner (spatial resolution 2.5 mm isotropic, 30 temporal frames; high-venc 150 cm/s; low-venc 30/50/70 cm/s). For retrospective tests, data were wrapped to venc values of 90/70/50/30 cm/s. To evaluate noise robustness, complex Gaussian noise was added at SNR levels of 20, 10, and 5 dB

(SNR = ∞ : no added noise). For prospective tests, unwrapping was performed on low-*venc* scans and ground truth was derived from dual-*venc* processing [13].

We compared PUDIP-Flow against LAP [9], NPRS [1], and PUMA [5]. Implementations of these baselines were taken from Dirix et al. [6].

2.4 Evaluation metrics

We evaluated unwrapping performance of PUDIP-Flow using the normalized root-mean-square error (nRMSE) over a manually segmented ROI. Let $\hat{\mathbf{v}}_n = [\hat{v}_{n,x}, \hat{v}_{n,y}, \hat{v}_{n,z}]^T$ and $\tilde{\mathbf{v}}_n = [\tilde{v}_{n,x}, \tilde{v}_{n,y}, \tilde{v}_{n,z}]^T$ denote the estimated and reference velocities at voxel n , respectively. The nRMSE is defined as

$$\text{nRMSE} = \sqrt{\frac{\sum_{n \in \text{ROI}} \|\hat{\mathbf{v}}_n - \tilde{\mathbf{v}}_n\|_2^2}{3 |\text{ROI}| (\tilde{v}_{\max} - \tilde{v}_{\min})^2}}, \quad (11)$$

where $|\text{ROI}|$ is the number of voxels in the ROI, and $\tilde{v}_{\max}$ and $\tilde{v}_{\min}$ are the maximum and minimum reference velocities within the ROI across the three encoding directions and all time frames. The factor 3 accounts for normalizing for the three velocity components.

Computation time was measured on the same case (data size $96 \times 104 \times 24$). The ROI mask lay within an $87 \times 32 \times 12$ bounding box and contained 6252 non-zero voxels.

3 Results

3.1 Retrospective experiments

Fig. 2 presents the unwrapping results of LAP, NPRS, PUMA, and PUDIP-Flow with manual segmentation mask or automated weighting mask on the wrapped flow data in the retrospective experiment under *venc* values of 90, 70, 50, 30 cm/s and SNR levels of ∞ , 20, 10, 5 dB. PUDIP-Flow achieves the lowest nRMSE in most settings and remains stable without a manual mask, unlike NPRS and PUMA.

3.2 Prospective experiments

Table 1 reports prospective results at *venc* = 70/50/30 cm/s. With a manual segmentation mask, PUDIP-Flow achieves the lowest nRMSE at 70 and 50 cm/s (NPRS is best at 30 cm/s); without a manual mask, PUDIP-Flow yields the lowest nRMSE across all *venc* values, and even yielded lower errors than NPRS with a manual mask.

Fig. 3 presents representative head-foot unwrapped velocity maps from three prospective datasets (*venc* = 70, 50, and 30 cm/s) without manual segmentation masks. NPRS and PUMA show strong mask dependence, with errors increasing




Fig. 2. The nRMSE of baselines and PUDIP-Flow across venc and SNR settings. (A): with segmentation mask; (B): without segmentation mask. ★ indicates that the method has the best performance.

Table 1. The nRMSE calculated based on the unwrapping results of LAP, NPRS, PUMA, and PUDIP-Flow under the vencs of 70, 50, and 30 cm/s in the prospective experiments.

Venc (cm/s)	nRMSE (%) w/ the segmentation mask			
	LAP	NPRS	PUMA	PUDIP-Flow
70	2.92 ± 1.34	2.06 ± 0.56	2.06 ± 0.52	1.43 ± 0.50
50	5.93 ± 0.75	3.61 ± 0.63	5.25 ± 5.07	3.55 ± 0.99
30	6.53 ± 0.76	5.06 ± 1.73	6.50 ± 3.74	5.23 ± 0.72
Time (s)	32.53	<u>41.94</u>	121.27	202.84
Venc (cm/s)	nRMSE (%) w/o the segmentation mask			
	LAP	NPRS	PUMA	PUDIP-Flow
70	2.92 ± 1.34	2.89 ± 0.98	8.05 ± 8.03	1.46 ± 0.49
50	5.93 ± 0.75	5.88 ± 0.95	11.92 ± 8.20	3.12 ± 0.64
30	6.53 ± 0.76	6.30 ± 0.96	6.93 ± 0.76	4.68 ± 0.86
Time (s)	32.53	<u>155.20</u>	4870.87	457.02

substantially when the mask is removed. In contrast, PUDIP-Flow remains robust and performs best when using the automated weighting mask. It has even higher performance than using the manual segmentation mask for low vencs.

Ablation studies (Table 2 and Table 3) indicate that the default PUDIP-Flow configuration is overall reasonable. Compared with PUDIP-Flow, changing the U-Net width ($N_c = 64$ or $N_c = 192$) or using the L2 loss generally results in dif-




Fig. 3. Representative head-foot velocity maps for LAP, NPRS, PUMA, and PUDIP-Flow without manual segmentation masks ($venc = 70/50/30$ cm/s).

Table 2. The nRMSE calculated based on the unwrapping results of different ablation studies under the $venc$ s of 70, 50, and 30 cm/s in the prospective experiments. $\uparrow$ indicates increased nRMSE relative to the default PUDIP-Flow.

Venc (cm/s)	nRMSE (%) w/ the segmentation mask				
	PUDIP-Flow	$N_c = 64$	$N_c = 192$	L2 Loss	Spatial TV
70	1.43 ± 0.50	1.36 ± 0.45	1.40 ± 0.51	1.53 ± 0.55	$1.94 \pm 0.84 \uparrow$
50	3.55 ± 0.99	$3.65 \pm 1.01 \uparrow$	3.53 ± 0.89	$3.67 \pm 0.81 \uparrow$	$3.67 \pm 0.78 \uparrow$
30	5.23 ± 0.72	5.22 ± 0.75	$5.30 \pm 0.73 \uparrow$	$5.36 \pm 0.74 \uparrow$	4.36 ± 0.96
Venc (cm/s)	nRMSE (%) w/o the segmentation mask				
	PUDIP-Flow	$N_c = 64$	$N_c = 192$	L2 Loss	Spatial TV
70	1.46 ± 0.49	1.42 ± 0.45	1.43 ± 0.48	$1.60 \pm 0.53 \uparrow$	$28.31 \pm 1.94 \uparrow$
50	3.12 ± 0.64	$3.21 \pm 0.69 \uparrow$	$3.18 \pm 0.68 \uparrow$	$3.51 \pm 0.77 \uparrow$	$26.21 \pm 3.23 \uparrow$
30	4.68 ± 0.86	4.71 ± 0.84	4.64 ± 0.89	$5.10 \pm 0.88 \uparrow$	$14.42 \pm 2.93 \uparrow$

ferent degrees of nRMSE increase across $venc$ values. Replacing spatial-temporal TV with spatial-only TV further degrades performance, especially without a segmentation mask.

We also compared different automated weighting masks, $\mathbf{P}_{std}$ (Eq. 8), $\mathbf{P}$ (Eq. 7), $\bar{\mathbf{P}}$ (temporal mean of Eq. 7), and an all-ones mask, in the setting without the

Table 3. The nRMSE calculated based on the unwrapping results using PUDIP-Flow with different automated weighting masks of the prospective experiments. $\uparrow$ indicates increased nRMSE relative to the default PUDIP-Flow.

Venc (cm/s)	nRMSE (%) w/o the segmentation mask			
	$\mathbf{P}_{\text{std}}$	$\mathbf{P}$	$\bar{\mathbf{P}}$	All ones
70	1.46 ± 0.49	1.25 ± 0.41	1.47 ± 0.69	$\uparrow 2.37 \pm 1.49$
50	3.12 ± 0.64	3.66 ± 1.04	$\uparrow 4.33 \pm 1.26$	$\uparrow 5.46 \pm 0.94$
30	4.68 ± 0.86	5.16 ± 0.85	$\uparrow 5.74 \pm 0.66$	$\uparrow 6.26 \pm 0.75$

manual segmentation mask. Overall, $\mathbf{P}_{\text{std}}$ is the most robust across venc values, whereas other masks tend to increase nRMSE.

4 Discussion

Both retrospective and prospective experiments demonstrated that PUDIP-Flow achieves the lowest nRMSE across nearly all wrapping levels and SNR conditions. A key advantage is its robustness without segmentation masks: while NPRS and PUMA exhibited substantial performance degradation without mask guidance, PUDIP-Flow maintained consistent accuracy using only the automated weighting mask. In the prospective experiments at $\text{venc} = 50$ and 30 cm/s, the automated mask even reduced errors compared with manual segmentation, likely because the soft, temporally varying weighting avoids hard boundary artifacts inherent in static manual masks.

The ablation study confirmed several design choices. The L1 loss outperformed L2, as L2 amplifies large errors in background air regions where phase values are unstable. Spatial-temporal TV proved critical: replacing it with spatial-only TV substantially degraded performance without segmentation masks, since the model could no longer exploit temporal continuity. Among automated masks, $\mathbf{P}_{\text{std}}$ was the most robust across venc values, as incorporating temporal standard deviation effectively highlights pulsatile vascular regions.

One limitation of PUDIP-Flow is its computation time (~ 8 min per case without ROI cropping), which could be mitigated by more efficient architectures. Additionally, validation was restricted to aortic flow; extension to other vascular territories may require adaptation of the mask design and loss function.

5 Conclusion

We developed PUDIP-Flow, an unsupervised, segmentation-free phase unwrapping method for aortic 4D flow MRI based on deep image prior. Unlike NPRS and PUMA, whose performance depends heavily on segmentation masks, PUDIP-Flow reliably unwraps phase using only an automatically computed weighting mask, achieving superior or comparable accuracy across a wide range of venc and

SNR settings. Its unsupervised and segmentation-free nature makes it a practical tool for improving the accuracy of flow quantification and VNR of clinical 4D flow MRI.

Acknowledgments. This study was funded in part by National Natural Science Foundation of China (Grant/Award Number: 82302295); Beijing Natural Science Foundation (Grant/Award Number: 1S24094).

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

References

1. Abdul-Rahman, H., Gdeisat, M., Burton, D., Lalor, M.: Fast three-dimensional phase-unwrapping algorithm based on sorting by reliability following a non-continuous path. In: Optical measurement systems for industrial inspection IV. vol. 5856, pp. 32–40. SPIE (2005)
2. Aristova, M., Pang, J., Ma, Y., Ma, L., Berhane, H., Rayz, V., Markl, M., Schnell, S.: Accelerated dual-vcnc 4d flow mri with variable high-vcnc spatial resolution for neurovascular applications. *Magnetic resonance in medicine* **88**(4), 1643–1658 (2022)
3. Berglund, J., Liljeblad, M., Baron, T.: Unwrapping phase contrast mri by iterative graph cuts. *Magnetic Resonance in Medicine* **92**(4), 1484–1495 (2024)
4. Berhane, H., Scott, M.B., Barker, A.J., McCarthy, P., Avery, R., Allen, B., Malaisrie, C., Robinson, J.D., Rigsby, C.K., Markl, M.: Deep learning-based velocity antialiasing of 4d-flow mri. *Magnetic resonance in medicine* **88**(1), 449–463 (2022)
5. Bioucas-Dias, J.M., Valadao, G.: Phase unwrapping via graph cuts. *IEEE Transactions on Image processing* **16**(3), 698–709 (2007)
6. Dirix, P., Jacobs, L., Kozerke, S.: 4dflow-unwrap: A collection of python-based phase unwrapping algorithms. *Intl Soc Mag Reson Med* (2024)
7. Ha, H., Kim, G.B., Kweon, J., Lee, S.J., Kim, Y.H., Lee, D.H., Yang, D.H., Kim, N.: Hemodynamic measurement using four-dimensional phase-contrast mri: quantification of hemodynamic parameters and clinical applications. *Korean journal of radiology* **17**(4), 445–462 (2016)
8. Itoh, K.: Analysis of the phase unwrapping algorithm. *Applied optics* **21**(14), 2470–2470 (1982)
9. Loecher, M., Schrauben, E., Johnson, K.M., Wieben, O.: Phase unwrapping in 4d mr flow with a 4d single-step laplacian algorithm. *Journal of Magnetic Resonance Imaging* **43**(4), 833–842 (2016)
10. Markl, M., Frydrychowicz, A., Kozerke, S., Hope, M., Wieben, O.: 4d flow mri. *Journal of Magnetic Resonance Imaging* **36**(5), 1015–1036 (2012)
11. Nett, E.J., Johnson, K.M., Frydrychowicz, A., Del Rio, A.M., Schrauben, E., Francois, C.J., Wieben, O.: Four-dimensional phase contrast mri with accelerated dual velocity encoding. *Journal of Magnetic Resonance Imaging* **35**(6), 1462–1471 (2012)
12. Pritt, M.D.: Congruence in least-squares phase unwrapping. In: IGARSS’97. 1997 IEEE International Geoscience and Remote Sensing Symposium Proceedings. Remote Sensing-A Scientific Vision for Sustainable Development. vol. 2, pp. 875–877. IEEE (1997)

13. Schnell, S., Ansari, S.A., Wu, C., Garcia, J., Murphy, I.G., Rahman, O.A., Rahsepar, A.A., Aristova, M., Collins, J.D., Carr, J.C., et al.: Accelerated dual-venc 4d flow mri for neurovascular applications. *Journal of Magnetic Resonance Imaging* **46**(1), 102–114 (2017)
14. Song, S.M.H., Napel, S., Pelc, N.J., Glover, G.H.: Phase unwrapping of mr phase images using poisson equation. *IEEE Transactions on Image Processing* **4**(5), 667–676 (1995)
15. Ulyanov, D., Vedaldi, A., Lempitsky, V.: Deep image prior. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 9446–9454 (2018)
16. Untenberger, M., Hüllebrand, M., Tautz, L., Joseph, A.A., Voit, D., Merboldt, K.D., Frahm, J.: Spatiotemporal phase unwrapping for real-time phase-contrast flow mri. *Magnetic resonance in medicine* **74**(4), 964–970 (2015)
17. Vishnevskiy, V., Walheim, J., Kozerke, S.: Deep variational network for rapid 4d flow mri reconstruction. *Nature Machine Intelligence* **2**(4), 228–235 (2020)
18. Yang, F., Pham, T.A., Brandenburg, N., Lütolf, M.P., Ma, J., Unser, M.: Robust phase unwrapping via deep image prior for quantitative phase imaging. *IEEE Transactions on Image Processing* **30**, 7025–7037 (2021)
19. Zhang, J., Rothenberger, S.M., Brindise, M.C., Scott, M.B., Berhane, H., Baraboo, J.J., Markl, M., Rayz, V.L., Vlachos, P.P.: Divergence-free constrained phase unwrapping and denoising for 4d flow mri using weighted least-squares. *IEEE transactions on medical imaging* **40**(12), 3389–3399 (2021)