



HAL
open science

Phase Unwrapping of Color Doppler Echocardiography using Deep Learning

Hang Jung Ling, Olivier Bernard, Nicolas Ducros, Damien Garcia

► **To cite this version:**

Hang Jung Ling, Olivier Bernard, Nicolas Ducros, Damien Garcia. Phase Unwrapping of Color Doppler Echocardiography using Deep Learning. IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control, 2023, 10.1109/TUFFC.2023.3289621 . hal-04142824v1

HAL Id: hal-04142824

<https://hal.science/hal-04142824v1>

Submitted on 27 Jun 2023 (v1), last revised 11 Jul 2023 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.




Distributed under a Creative Commons Attribution 4.0 International License

Phase Unwrapping of Color Doppler Echocardiography using Deep Learning

Hang Jung Ling, Olivier Bernard, Nicolas Ducros, and Damien Garcia

Abstract—Color Doppler echocardiography is a widely used non-invasive imaging modality that provides real-time information about the intracardiac blood flow. In an apical long-axis view of the left ventricle, color Doppler is subject to phase wrapping, or aliasing, especially during cardiac filling and ejection. When setting up quantitative methods based on color Doppler, it is necessary to correct this wrapping artifact. We developed an unfolded primal-dual network to unwrap (dealias) color Doppler echocardiographic images and compared its effectiveness against two state-of-the-art segmentation approaches based on nnU-Net and transformer models. We trained and evaluated the performance of each method on an in-house dataset and found that the nnU-Net-based method provided the best dealiased results, followed by the primal-dual approach and the transformer-based technique. Noteworthy, the primal-dual network, which had significantly fewer trainable parameters, performed competitively with respect to the other two methods, demonstrating the high potential of deep unfolding methods. Our results suggest that deep learning-based methods can effectively remove aliasing artifacts in color Doppler echocardiographic images, outperforming DeAN, a state-of-the-art semi-automatic technique. Overall, our results show that deep learning-based methods have the potential to effectively preprocess color Doppler images for downstream quantitative analysis.



Index Terms—Color Doppler, Flow imaging, Echocardiography, Phase unwrapping, Dealiasing, Deep learning, Deep unfolding, U-Net, Transformer, Primal-dual.

I. INTRODUCTION

COLOR Doppler ultrasound is a widely accepted clinical imaging modality for non-invasive, real-time analysis of cardiovascular blood flow. While two-dimensional color Doppler is commonly used for qualitative mapping of flow characteristics, its applications for quantitative analysis are limited. Common cardiovascular applications of color Doppler include detection of valvular diseases [1] and septal defects [2], or guiding the positioning of the pulsed-wave sample volume for spectral Doppler [3]. Among methods to make color Doppler quantitative, Vector Flow Mapping has been introduced for intracardiac flow dynamics. This method allows the computation of 2D or 3D intraventricular velocity vector maps from color Doppler fields, using a physically constrained optimization approach [4], [5]. Intracardiac vector flow mapping from color Doppler requires two prerequisite steps: 1) delineation of the endocardial inner wall, and 2) correction of wrapped (aliased) Doppler regions.

With respect to the first step, Painchaud *et al.* [6] recently introduced a 2D+time deep learning architecture to enforce temporal consistency and smoothness from one frame to the next. The second step of correction is necessary due to the

occurrence of aliasing, which is an artifact resulting from insufficient slow-time sampling. This issue arises when the pulse repetition frequency (PRF) is unable to capture high axial velocities effectively. This causes the Doppler velocity to be wrapped to the opposite side of the Doppler spectrum when its absolute value exceeds the Nyquist velocity. Experienced clinicians can easily identify zones of aliasing in most color Doppler images, where the color-coded velocities shift from red to blue or vice versa. Aliasing can be removed in Doppler echocardiography by designing multi-PRF sequences, as described by Posada *et al.* [7]. However, this approach requires control of the ultrasound machine and is primarily suitable for high-frame-rate echocardiography. When clinical scanners are used, aliasing must be corrected by post-processing the color Doppler fields. While a number of unwrapping algorithms have been proposed for dealiasing data maps in atmospheric science, geodesy, and optical interferometry [8]–[10], this problem has received less attention in color Doppler imaging.

Inspired by traditional radar approaches, Muth *et al.* [11] developed a segmentation-based method for color Doppler dealiasing using statistical region merging, called DeAN. This unsupervised method uses a scalar hyperparameter to control the segmentation process. An optimal parameter was determined from a supervised analysis of 50 color Doppler data. However, it turns out that the DeAN method fails in difficult cases as shown in Fig. 11 in [11] and that supervised corrections are still necessary in some situations. With the goal

Hang Jung Ling, Olivier Bernard, Nicolas Ducros and Damien Garcia are with CREATIS, CNRS UMR 5220, INSERM U1294, INSA-Lyon, University of Lyon 1, 69100 Villeurbanne, France (e-mail: hang-jung.ling@creatis.insa-lyon.fr; damien.garcia@creatis.insa-lyon.fr). Nicolas Ducros is also with Institut Universitaire de France (IUF).

Highlights

- Our deep-unfolding-based primal-dual network (PDNet) incorporated the forward operator as prior information and had only 0.03M parameters.
- Our deep learning (DL) models outperformed a state-of-the-art non-DL approach in phase unwrapping of color Doppler echocardiography, with nnU-Net being the best candidate, followed by PDNet with 233 times fewer parameters.
- Automatic and accurate color Doppler echocardiographic phase unwrapping ensures correct visualization and enables quantification of intracardiac blood flow.

of developing imaging tools that quantify blood flow from color Doppler, we propose a deep learning (DL) approach to correct the aliased areas of echocardiographic color Doppler maps. DL has been proposed for color Doppler dealiasing in vascular flow imaging by Nahas *et al.* [12]. Their approach aimed to solve the double aliasing problem using two U-Nets. The first U-Net detected the presence of single aliasing while the second U-Net was trained to identify and segment double-aliased pixels. They evaluated the performance of their model by training it with different types of ultrasound information. They found that the model trained with a combination of Doppler frequency, power, and bandwidth performed the best for dealiasing in the femoral bifurcation.

In our work, we focused on Doppler echocardiography. In contrast to vascular flow imaging, cardiac color flow imaging can be subject to substantial clutter signals originating from the myocardium and tending to spread the aliased patterns. With the goal of proposing a robust DL method that correctly handles aliasing in most situations, we developed and compared several architectures. Our main contributions are:

- 1) We designed a primal-dual network based on the idea of deep unfolding, and compared it with state-of-the-art DL segmentation methods and DeAN.
- 2) We used a private color Doppler echocardiographic dataset acquired in apical three-chamber view (45 patients, 1,338 aliased and 2,379 non-aliased frames) to train the neural networks and analyze their performance.
- 3) We investigated the value of adding Doppler power as input information to improve dealiasing.
- 4) We introduced a data augmentation strategy that generates synthetic aliasing, which solved the class imbalance problem and improved dealiasing performance on difficult color Doppler images.

II. METHODOLOGY

Aliasing artifacts occur when axial blood speeds (velocity magnitudes) exceed the Nyquist velocity V_N . The acquired Doppler velocity V_D can be written as a function of the unwrapped or alias-free Doppler velocity V_u as follows:

$$V_D = V_u - 2 \times n_N V_N, \quad (1)$$

where n_N is an integer called the Nyquist number, which represents the number of times the signal wraps around the Nyquist limit. The Nyquist number reads (see [7] for the demonstration)

$$n_N = \text{floor} \left(\frac{V_u + V_N}{2V_N} \right). \quad (2)$$

Except for highly turbulent flows that may occur in transvalvular or transseptal jets, there is no multiple aliasing in the adult left ventricle scanned in the apical long-axis view, *i.e.*, the integer n_N belongs to $\{-1, 0, 1\}$. Indeed, in adult echocardiography with a 3 MHz phased array, Nyquist velocities typically range from 0.55 to 0.7 m/s. Thus, single (*i.e.*, $n_N = -1$ or 1) or no (*i.e.*, $n_N = 0$) wrapping occurs as long as the actual blood speed is less than 1.65-2.1 m/s (see (4)). It follows that double aliasing does not occur in the left ventricle in most situations without valvular disease or cardiac shunt. Equations (1) and (2) can be rewritten to express V_D as a wrapped version of V_u :

$$V_D = K(V_u) = (V_u + V_N) \bmod (2V_N) - V_N, \quad (3)$$

where \bmod is the modulo operation. In particular, for $n_N \in \{-1, 0, 1\}$, the wrapping function K becomes

$$V_D = K(V_u) = \begin{cases} V_u - 2V_N & \text{if } V_N < V_u < 3V_N \\ V_u & \text{if } -V_N \leq V_u \leq V_N \\ V_u + 2V_N & \text{if } -3V_N < V_u < -V_N \end{cases}. \quad (4)$$

This representation implies that the dealiasing problem can be approached in two different ways: *i)* a problem that inverts the wrapping function (3) and recovers V_u from the Doppler velocities V_D by changing absolute jumps greater than V_N to their $2 \times V_N$ complement; *ii)* a multi-class segmentation approach that assigns a Nyquist number n_N (2) to each pixel of the input image then computes the actual unwrapped velocities using (1). We investigated three deep learning (DL) models for dealiasing color Doppler. The first method was derived from the deep unfolding/unrolling framework and solved the inverse problem defined in (3) to estimate the actual velocities. We faced a nonlinear inverse problem on non-trivial data, whose solution may contain phase jumps at the blood/myocardium interfaces. Unrolled methods are well suited for solving inverse problems. Primal-dual optimization, on the other hand, is useful for nonlinear problems. For these reasons, we tested a learned primal-dual algorithm inspired by Adler *et al.* [13], as described in the following subsection. The other two methods were state-of-the-art networks that we have adapted for determining Nyquist numbers in color Doppler images. Fig. 1 illustrates the pipeline used for all three methods, whose input data were the Doppler velocity multiplied by the Doppler power before scan-conversion.

A. OriPDNet: A Primal-Dual-based Deep Unfolding Network to Solve Inverse Problems

To solve our nonlinear inverse problem (3), we used OriPDNet (refer to Fig. 2 in [13] for the network architecture),

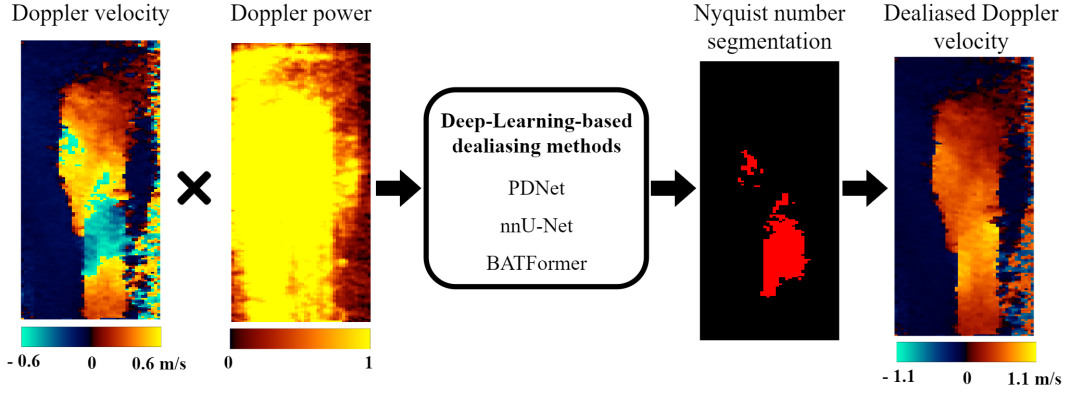


Fig. 1: Pipeline of the deep learning-based methods for color Doppler dealiasing.

a deep unfolding network based on a primal-dual optimization scheme [13]. Given a general inverse problem aiming to obtain the solution f from the measurement g with the forward operator K :

$$g = K(f), \quad (5)$$

the outline of OriPDNet to solve this problem is presented in Algorithm 1.

Algorithm 1 OriPDNet: Original primal-dual network

Initialize $f_0, h_0 = [\mathbf{0}, \mathbf{0}, \mathbf{0}, \mathbf{0}, \mathbf{0}] \in \mathbb{R}^{M \times N \times 5}$
for $i = 1, \dots, I$ **do**
 $h_i \leftarrow \Gamma_{\theta_i^d}(h_{i-1}, K(f_{i-1}^{(2)}), g)$
 $f_i \leftarrow \Lambda_{\theta_i^p}(f_{i-1}, [\partial K(f_{i-1}^{(1)})]^*(h_i^{(1)}))$
end for
return $f_I^{(1)}$

OriPDNet involves several variables and operators, including the forward operator K , the adjoint of its Fréchet derivative $[\partial K]^*$, the input measured data g , the primal and dual variables f_i and h_i , and the learned primal and dual proximal operators $\Lambda_{\theta_i^p}$ and $\Gamma_{\theta_i^d}$. Convolutional layers are used to learn these proximal operators. The hyperparameter I , which determines the number of iterations, requires careful tuning for each specific problem. The primal and dual variables, f_i and h_i , are initialized then iteratively updated using the learned primal and dual proximal operators, $\Lambda_{\theta_i^p}$ and $\Gamma_{\theta_i^d}$. The solution to the inverse problem (5) is obtained by extracting the first element of the primal variables, $f_I^{(1)}$. In [13], the authors recommended setting the dimension of the primal and dual spaces to five as the best compromise between memory usage and reconstruction quality, *i.e.*, $f_i = [f_i^{(1)}, f_i^{(2)}, f_i^{(3)}, f_i^{(4)}, f_i^{(5)}] \in \mathbb{R}^{M \times N \times 5}$ and $h_i = [h_i^{(1)}, h_i^{(2)}, h_i^{(3)}, h_i^{(4)}, h_i^{(5)}] \in \mathbb{R}^{M \times N \times 5}$, where $(M \times N)$ is the size of the input data. We conducted preliminary testing and validated the use of five spaces, in accordance with their suggestion.

B. PDNet: A Deep Unfolding Network for Color Doppler Dealiasing

To deal with our specific inverse problem (3) for color Doppler dealiasing, we adapted OriPDNet. The modified ver-

sion was named as PDNet and summarized in Algorithm 2, with the main changes highlighted in blue.

Algorithm 2 PDNet: Proposed primal-dual network

Initialize $V_0, h_0 = [\mathbf{0}, \mathbf{0}, \mathbf{0}, \mathbf{0}, \mathbf{0}] \in \mathbb{R}^{M \times N \times 5}$
for $i = 1, \dots, I$ **do**
 $h_i \leftarrow \Gamma_{\theta_i^d}(h_{i-1}, K(V_{i-1}^{(2)}), V_D)$
 $V_i \leftarrow \Lambda_{\theta_i^p}(V_{i-1}, h_i^{(1)})$
end for
return $C_\theta(V_I^{(1)})$

with $V_u = V_I^{(1)}$ and $n_N = C_\theta(V_I^{(1)})$.

Specifically, we defined the forward operator K as a wrapping function given by (3). Despite the discontinuity of this function at each $V = V_N \pm 2kV_N$ (with $k \in \mathbb{Z}^*$), its derivative was an identity function, *i.e.*, $\partial K(V) = \text{id}$. Thus, its adjoint $[\partial K(V)]^*$ was also an identity function. Unlike the original approach (see Algorithm 1), we used the same feature maps for each iteration of the main loop, which significantly reduced the number of parameters to learn (30,000 instead of $30,000 \times I$, with I being the number of iterations) while maintaining the same accuracy. We made this change to avoid training instabilities that we observed while experimenting with OriPDNet. We also added a convolutional layer C_θ at the end of the network to output the Nyquist number from the estimated velocities $V_I^{(1)}$. The main reason for this was to avoid non-integer Nyquist numbers due to the regressed velocities. For a fair comparison between PDNet and OriPDNet, the same convolution layer C_θ was also applied to the output of OriPDNet.

C. Segmentation Networks for Color Doppler Dealiasing

nnU-Net is currently one of the best performing approaches for medical image segmentation [15]. This model is based on the U-Net architecture and implements several successful DL tricks, such as automatic hyperparameter search of the U-Net architecture to increase accuracy, a patch-wise approach to preserve image resolution, a deep supervision strategy to maintain accuracy at all scales, and data augmentation during both training and inference to enforce generalization. In this

TABLE I: Main configurations of the three methods evaluated in this study. *Lowest resolution:* Size of the lowest resolution of feature maps in pixels. *Down. scheme:* Downsampling scheme. *Up. scheme:* Upsampling scheme. *Optimization scheme:* Optimizer + initial learning rate (+ learning rate scheduler used). *# param.:* Number of trainable parameters.

Methods	Number of feature maps	Lowest resolution	Down. scheme	Up. scheme	Normalization scheme	Batch size	Optimization scheme	Loss function	# param.
PDNet	32 → 32 → 5	192×40	-	-	-	4	ADAM [14] + 0.001 + cosine annealing	Cross entropy + Dice	0.03M
nnU-Net	32 ↓ 480 ↑ 32	12×5	Stride pooling	Deconv.	Instance norm.	4	SGD + 0.01 + polynomial decay	Cross entropy + Dice	7M
BATFormer	16 ↓ 256 ↑ 16	16×16	Max pooling	2 × 2 repeats	Batch norm.	4	ADAM + 0.001	Cross entropy + Dice + smooth L1	1.2M

study, we addressed the dealiasing of color Doppler as a 3-class segmentation problem with nnU-Net, where each class corresponded to a Nyquist number $n_N \in \{-1, 0, 1\}$. Our network included four stages in the encoder/decoder parts and had an input size of 192×40 pixels, which was the median image size of our dataset. Table I provides more details about the architecture and the training scheme (see Fig. C.1 in [15] for an illustration of the nnU-Net architecture).

Recently, transformer-based approaches have been shown to outperform the nnU-Net model in some medical challenges [16]. These models attempt to solve the segmentation problem in a different way, using attention mechanisms with receptive fields that cover the entire image. Among the best-performing models, we chose to train the BATFormer architecture [17] for the color Doppler dealiasing task, using a segmentation-based technique. This model employs a multiscale approach based on a U-Net architecture, with transformer blocks added to the decoder. This strategy results in an efficient and lightweight architecture (1.2M parameters) that is suitable for learning from small to medium-sized datasets. The main configurations of BATFormer are listed in Table I, and an illustration of its architecture can be found in Fig. 2 in [17].

D. Input Data Strategy

Color Doppler echocardiography produces two types of information: *i*) Doppler velocity, which can be corrupted by aliasing in regions of high blood speed, and *ii*) Doppler power, which provides insight into the regions where velocity measurements are reliable. Using both Doppler power and velocity as input to DL models allowed them to learn how to limit the dealiasing process in regions of interest and identify ambiguous areas. Therefore, we performed an ablation study to evaluate the potential improvement provided by Doppler power. This study was conducted using the nnU-Net architecture, known for its stability in training and optimal configurations. Specifically, three nnU-Net models were trained with three combinations of input data: 1) nnU-Net #1 trained with Doppler velocity only, 2) nnU-Net #2 trained with the concatenation of Doppler velocity with Doppler power, and 3) nnU-Net #3 trained with the multiplication of Doppler velocity by Doppler power. After determining the best candidate for the input data, we trained the three models, PDNet, nnU-Net, and BATFormer, using this input combination to compute the

Nyquist numbers n_N (2) from which the unwrapped Doppler velocities V_u (1) were derived. Our goal was to increment or decrement the Doppler velocities by $2n_N V_N$, not to modify them by smoothing, for example.

E. Artificial Aliasing Augmentation Strategy

Color Doppler images may exhibit aliasing only in localized regions or frames, resulting in datasets that are often imbalanced, with most pixels belonging to the background class (*i.e.*, without aliasing). To address this issue, we used standard data augmentation techniques such as rotation, flipping etc., during training. We also proposed an additional data augmentation technique, which we called artificial aliasing augmentation, to improve the generalizability of our algorithms. This technique involved identifying regions with high Doppler velocity and power on alias-free Doppler images, and applying a wrapping function defined by (3) with a lower Nyquist velocity to create artificial aliasing artifacts, followed by a normalization. The ground-truth references of these artificially aliased frames were created on the fly by comparing the Doppler velocities before and after this augmentation. By creating realistic artificially aliased frames, as shown in Fig. 2, this strategy enabled us to balance the classes in training batches. To evaluate the potential benefits of artificial aliasing augmentation, we conducted an additional ablation study where we tested the three DL models with and without this technique during training.

F. DeAN: State-of-the-Art non DL-Based Dealiasing Method for CDI

DeAN, mentioned in Section I, is currently one of the most powerful non DL-based methods for color Doppler dealiasing. It is a semi-supervised method with a hyperparameter, Q . To unwrap aliased pixels, DeAN first segments color Doppler images using a region-merging scheme based on the Hoeffding's probability inequality. Then, DeAN compares each segmented region with its nearest neighbors and performs dealiasing if necessary. This step is based on the assumption that the largest segment is not aliased and is repeated until all the segments have been analyzed. The main drawback of this method is the need to manually search for the optimal Q hyperparameter for each frame to obtain the best dealiasing results. We compared




Fig. 2: Generation of synthetic images with artificial aliasing artifacts (right column) from non-aliased images (left column).

the dealiasing performance of DeAN with both the default $Q = 10$ and with the manually optimized Q hyperparameter, against the three DL methods.

G. Evaluation metrics

All three DL models were designed to output the Nyquist numbers n_N ; the dealiased velocity maps, V_u , were recovered using (1). To evaluate the accuracy of the dealiased velocity maps and the Nyquist numbers outputted by each method, we computed four evaluation metrics.

We compared the dealiased Doppler velocity maps V_u with the ground-truth alias-free Doppler velocity maps V_{ref} by computing the cosine similarity index:

$$\text{Cosim}(V_u, V_{ref}) = \frac{V_u \cdot V_{ref}}{\|V_u\| \cdot \|V_{ref}\|} \quad (6)$$

Cosine similarity is a commonly used similarity measure for comparing text data or images. We used this similarity index in our previous work on color Doppler dealiasing [11]. In addition, we computed three classification metrics to verify whether each pixel was classified correctly on color Doppler images. The first classification metric was the balanced accuracy score, which is more suitable for unbalanced datasets than the classical accuracy score. It was calculated using the following formula:

$$\text{Accuracy} = \frac{1}{2} \times \left(\frac{\text{TP}}{\text{TP} + \text{FN}} + \frac{\text{TN}}{\text{TN} + \text{FP}} \right) \quad (7)$$

where TP, FN, TN and FP refer to true positives, false negatives, true negatives, and false positives, respectively. Besides, the classical recall ($\frac{\text{TP}}{\text{TP} + \text{FN}}$) and precision ($\frac{\text{TP}}{\text{TP} + \text{FP}}$) metrics were also computed to evaluate the overall performance of the methods.

To ensure the reliability and relevance of the results, we conducted a 9-fold cross-validation to compute the scores presented in each table in Section III-B. For each fold, we split the dataset into training, validation, and test sets using a ratio of 36/4/5 patients. This resulted in an average of 2,974, 330, and 413 color Doppler frames for the training, validation, and test sets, respectively.

III. EXPERIMENT SETUP AND RESULTS

A. Dataset and Training Strategies

1) *Color Doppler dataset*: To evaluate the performance of our methods, we used a color Doppler echocardiographic dataset of 45 patients that were acquired using a Vivid 7 ultrasound system (GE Healthcare, USA) and a GE 5S cardiac sector probe (bandwidth = 2-5 MHz). Doppler velocity and power data prior to scan conversion were extracted into HDF formats using EchoPAC software (GE Healthcare). The EchoPAC software returned unitless power data in the range of 1-100. The power data, P , were compressed by taking the logarithm and then scaled to $[0,1]: \log(P)/2 \leftarrow P$. The cardiologist used the default settings (center frequency, pulse length, pulse repetition frequency, packet size, clutter filter, etc.). These proprietary parameters are masked and could not be extracted. In most cases, the Nyquist velocity was about 0.6 m/s. Assuming a center frequency of 3 MHz, the pulse repetition frequency was approximately 4500 Hz. These anonymized data came from a previous published study [18]. The patients were examined in the echocardiography laboratory under standard medical conditions. As a result, some patients had significant heart disease, while others had no visible pathology. The random selection of patients and their anonymization prevent us from knowing their demographic and pathological status. The sequences were acquired in the apical three-chamber view and included both Doppler velocity and power information. Each sequence covered at least one complete cardiac cycle, resulting in a total of 1,338 aliased frames and 2,379 non-aliased frames. Since color Doppler has a relatively low frame rate of 10 to 15 frames per second in clinical echocardiography, we considered each frame to be independent. To avoid interpolation artifacts, the data were collected and processed in polar coordinates (i.e., before scan-conversion), but for better visualization, the results were presented in Cartesian coordinates.

The training, validation, and test data sets, i.e., the pairs of original and alias-free Doppler velocity maps, were generated by an experienced analyzer. For this task, the non-scan-converted clinical Doppler maps were oversegmented and labeled using a statistical growing region method (see Fig. 1.B in [11]). The analyst manually identified the aliased regions, specifically focusing on those related to intraventricular blood flow, which were then corrected by applying $\pm 2V_N$. The noisy regions associated with low Doppler power were left unchanged. Examples of color Doppler frames along with their reference segmentations are illustrated in Fig. 3.

2) *Training strategies*: To train the deep learning (DL) methods described in Section II, we performed supervised learning using the ground-truth segmentations from our in-house dataset. Besides applying the data augmentation strategies mentioned in Section II-E, we further addressed the class imbalance of our dataset by ensuring that each batch contained at least one aliased image, whether real or synthetic. The BATFormer model was designed using the official implementation proposed in its GitHub repository¹. This model

¹<https://github.com/xianlin7/BATFormer>




Fig. 3: Generation of the ground-truth masks from the Doppler velocities. The red, black, and green segmentation masks correspond to Nyquist numbers $n_N = 1, 0,$ and $-1,$ respectively.

took color Doppler images resized to 256×256 pixels as input and was trained for 400 epochs. On the other hand, nnU-Net and PDNet were implemented using the ASCENT framework². For these two approaches, we used a patch-wise approach to preserve the resolution of the input data. The models were trained for 1000 epochs to prevent any potential under/overfitting. More details on the BATFormer, nnU-Net, and PDNet architectures are provided in Table I.

B. Experimental Results

1) *Doppler power information was useful in dealiasing difficult case:* Table II reports the results of the ablation study aimed at identifying the optimal combination of input data. The results indicate that the three nnU-Net models performed similarly across all metrics, implying that incorporating Doppler power information in the input data did not substantially improve the models' performance. However, upon evaluation on a challenging test set (right part of Table II), the model that was trained with the multiplication of Doppler velocity and Doppler power (nnU-Net #3) demonstrated better performance for all metrics except precision. The last two columns of Fig. 4 show two samples taken from the difficult fold, where the aliased and non-aliased regions had similar hues. This made the correction of the aliased velocities difficult. Thus, although not critical, using the Doppler velocity-Doppler power product as input data is recommended as it can enhance the models' generalization ability, especially for challenging data. For subsequent experiments and results, we trained all DL methods with this input combination.

2) *PDNet outperformed its original counterpart:* We conducted a study to determine the optimal number of iterations for updating the primal and dual variables in both OriPDNet and PDNet, given the sensitivity of this type of method to this parameter. The results are shown in Table III. From this table, we can see that OriPDNet reached a plateau after 10 iterations, and beyond 20 iterations, it became unstable during training and failed to produce results. On the other hand, PDNet exhibited greater training stability and reached a performance plateau after 20 iterations. Moreover, PDNet achieved better optimal results than OriPDNet for all metrics

except for the cosine similarity index, where both methods performed equally well. These results suggest that using the same feature maps in the primal-dual approach is more suitable for the dealiasing task and support the use of our deep unfolding network. Additionally, it is worth mentioning that PDNet had only 30,000 parameters, making it the lightest of the three DL models tested, as detailed in Table I. The significant reduction in parameters of PDNet as compared with other models was due to the inclusion of the forward operator as prior information and the use of the same feature maps per iteration.

3) *Artificial aliasing augmentation improved the performance of segmentation-based networks:* The results presented in Table IV show that the use of artificial aliasing augmentation during training had varying effects on the performance of the different DL models. For nnU-Net, there was a slight improvement in all metrics except precision. In contrast, BATFormer showed significant improvement in accuracy, recall, and precision metrics, with values increasing from 0.88, 0.76, and 0.85 to 0.91, 0.81, and 0.91, respectively. However, for PDNet, the use of artificial aliasing augmentation resulted in degraded performance, with accuracy and recall decreasing from 0.94 and 0.87 to 0.88 and 0.77, respectively. These results highlight the challenge that the primal-dual-based regression methods face when generalizing to different types of aliasing. Based on the findings shown in Tables II to IV, we determined that the multiplication of Doppler velocity and power should be used as input for all methods, while artificial aliasing augmentation should be applied during training only for the segmentation-based techniques, *i.e.*, nnU-Net and BATFormer.

4) *nnU-Net gave the best dealiasing results:* Table V presents the final results of our study, where we compare the three DL methods with their optimal configurations against the DeAN algorithm. We observe that all three DL methods outperformed the DeAN algorithm, even the version with the manually chosen optimal Q hyperparameter. This outcome confirms the potential of DL methods for color Doppler dealiasing. Among the DL methods, nnU-Net achieved the highest scores overall, with a cosine similarity close to 1, an accuracy of 0.96, a recall of 0.91, and a precision of 0.89. Therefore, we conclude that nnU-Net is the best DL approach currently available for dealiasing tasks in echocardiography.

²<https://github.com/creatis-myriad/ASCENT>

TABLE II: Dealiasing by nnU-Net trained with different combinations of input data using 9-fold cross-validation. #1, #2, and #3 correspond to the use of Doppler velocity only, the concatenation of Doppler velocity with Doppler power, and the multiplication of the Doppler velocity by Doppler power as input data, respectively. The *Difficult fold* column shows the evaluation results of different nnU-Net models on a challenging test set containing color Doppler frames with aliased and non-aliased regions of similar hue.

Methods	Full dataset (Number of frames = 3,717)				Difficult fold (Number of frames = 413)			
	Cosim $\pm\sigma$	Accuracy $\pm\sigma$	Recall $\pm\sigma$	Precision $\pm\sigma$	Cosim $\pm\sigma$	Accuracy $\pm\sigma$	Recall $\pm\sigma$	Precision $\pm\sigma$
nnU-Net #1 <i>velocity only</i>	0.99 ± 0.01	0.95 ± 0.03	0.90 ± 0.07	0.92 ± 0.06	0.98 ± 0.01	0.93 ± 0.06	0.86 ± 0.13	0.95 ± 0.01
nnU-Net #2 <i>velocity-power concatenation</i>	0.99 ± 0.01	0.95 ± 0.03	0.91 ± 0.07	0.91 ± 0.05	0.98 ± 0.01	0.94 ± 0.07	0.88 ± 0.13	0.92 ± 0.05
nnU-Net #3 <i>velocity-power multiplication</i>	0.99 ± 0.01	0.95 ± 0.06	0.89 ± 0.13	0.90 ± 0.11	0.99 ± 0.01	0.94 ± 0.07	0.89 ± 0.14	0.92 ± 0.02

TABLE III: 9-fold cross-validation dealiasing results of PDNet trained with different number of iterations (*# iter.*) for updating the primal and dual variables. The results on the left were obtained with OriPDNet, *i.e.*, different feature maps per iteration, while the results on the right were obtained with the proposed PDNet using the same feature maps for each iteration.

# iter.	OriPDNet: Different feature maps per iteration				PDNet: Same feature maps for each iteration			
	Cosim $\pm\sigma$	Accuracy $\pm\sigma$	Recall $\pm\sigma$	Precision $\pm\sigma$	Cosim $\pm\sigma$	Accuracy $\pm\sigma$	Recall $\pm\sigma$	Precision $\pm\sigma$
1	0.95 ± 0.03	0.86 ± 0.06	0.71 ± 0.13	0.51 ± 0.09	0.95 ± 0.03	0.87 ± 0.06	0.74 ± 0.12	0.52 ± 0.10
10	0.98 ± 0.02	0.92 ± 0.06	0.84 ± 0.13	0.78 ± 0.11	0.98 ± 0.02	0.92 ± 0.06	0.84 ± 0.12	0.80 ± 0.10
20	0.98 ± 0.02	0.92 ± 0.06	0.84 ± 0.12	0.80 ± 0.11	0.98 ± 0.01	0.94 ± 0.06	0.87 ± 0.13	0.83 ± 0.11
30	-	-	-	-	0.97 ± 0.02	0.91 ± 0.07	0.82 ± 0.13	0.77 ± 0.14

TABLE IV: 9-fold cross-validation dealiasing results of the three implemented deep learning solutions trained with and without the proposed artificial aliasing augmentation strategy.

Methods	Without artificial aliasing augmentation				With artificial aliasing augmentation			
	Cosim $\pm\sigma$	Accuracy $\pm\sigma$	Recall $\pm\sigma$	Precision $\pm\sigma$	Cosim $\pm\sigma$	Accuracy $\pm\sigma$	Recall $\pm\sigma$	Precision $\pm\sigma$
PDNet	0.98 ± 0.02	0.94 ± 0.06	0.87 ± 0.12	0.83 ± 0.10	0.98 ± 0.01	0.88 ± 0.08	0.77 ± 0.16	0.84 ± 0.09
nnU-Net	0.99 ± 0.01	0.95 ± 0.03	0.89 ± 0.09	0.90 ± 0.06	0.99 ± 0.01	0.96 ± 0.03	0.91 ± 0.06	0.89 ± 0.06
BATFormer	0.98 ± 0.02	0.88 ± 0.08	0.76 ± 0.17	0.85 ± 0.11	0.98 ± 0.02	0.91 ± 0.07	0.81 ± 0.14	0.91 ± 0.07

Additionally, it is interesting to note that BATFormer showed a clear improvement when we increased the amount of synthetic data, indicating that this type of approach requires a larger dataset to improve its performance for the dealiasing task. Finally, it is worth noting that PDNet achieved promising results with 233 times fewer parameters compared to nnU-Net, highlighting the potential of incorporating analytical context into the DL framework to regularize the solution space.

We also provide a visual inspection of the performance of the various methods on aliased images with different degrees of difficulty in Fig. 4. We can see that the DL methods

performed similarly well on the easy and moderate cases (first two columns), but nnU-Net produced the closest results to the reference on the more challenging case (third column). This finding is consistent with the quantitative results presented in Table V. The last column in Fig. 4 shows an example where no method was able to handle aliasing correctly. This example is similar to the one in the first column, but with a more pronounced level of aliasing. In this particular case, the DeAN method gave the best results. This suggests that it would be advisable to supplement the training set with challenging configurations.

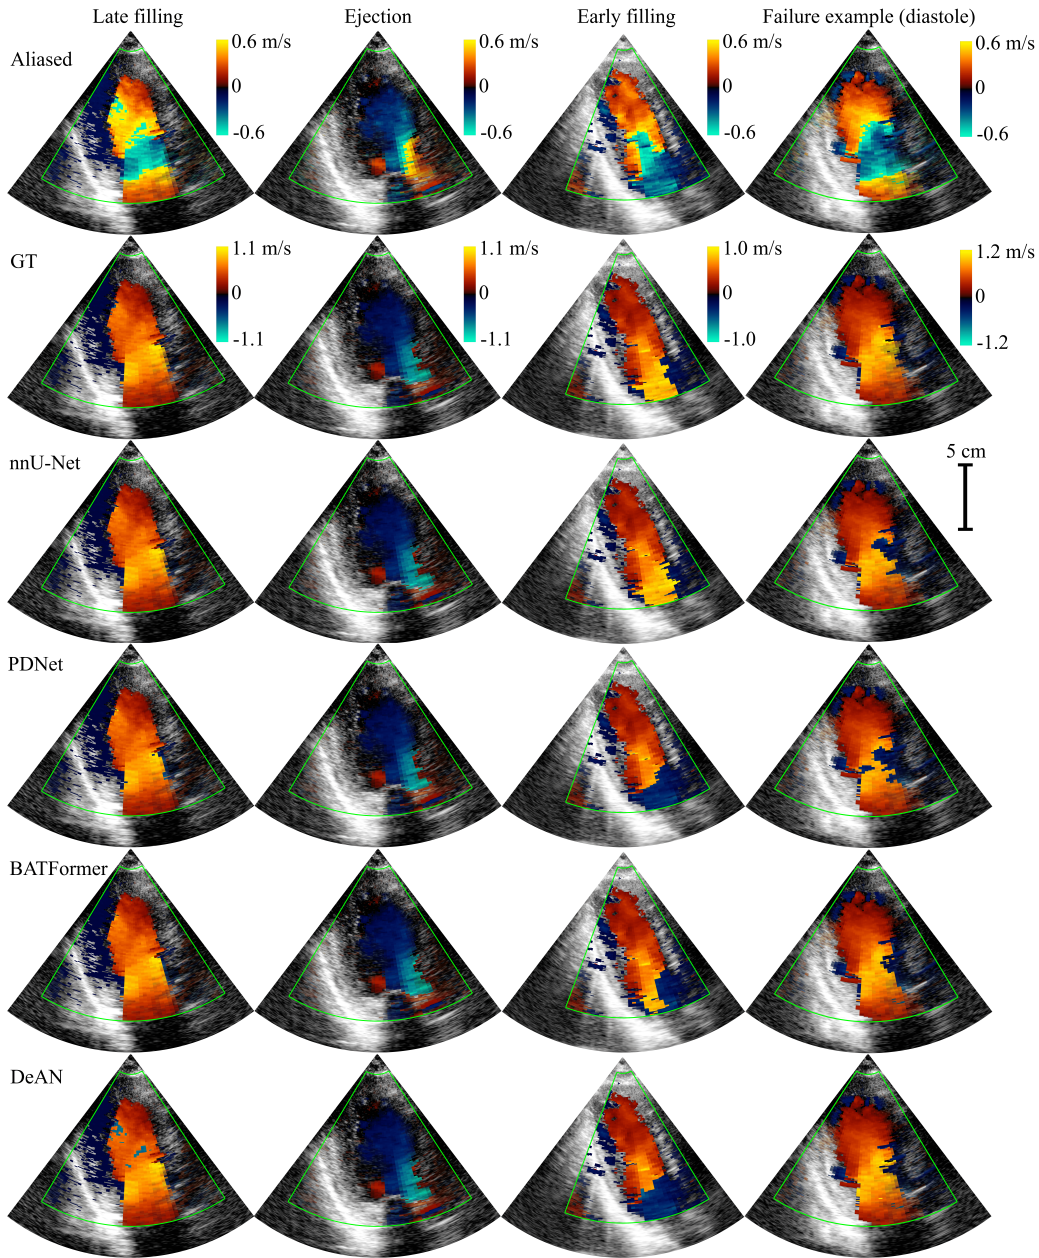


Fig. 4: Color Doppler images acquired during: from left to right column, late diastolic filling, systolic ejection, early diastolic filling, and diastole (a failed case). They were dealiased by DL-based methods and by DeAN with optimized Q hyperparameters. First row: aliased raw color Doppler. Second row: alias-free ground truth (GT).

IV. DISCUSSIONS

Color Doppler imaging takes high-pass-filtered in-phase/quadrature (I/Q) data of the same region of interest acquired along the slow-time axis and differentiates them pairwise using a lag-1 autocorrelator. The resulting maps show blood displacements between two consecutive slow-time samples. By its very nature, color Doppler imaging is an interferometric technique that enables the measurement of displacements with a precision that can reach fractions of the center wavelength. Similarly, synthetic aperture radar interferometry (InSAR), a remote sensing technique used to map the Earth's surface deformations, generates interferograms that display ground-surface displacements.

Like color Doppler, most interferometric imaging techniques in fields such as medical imaging, remote sensing, and optical metrology (e.g., phase-contrast MRI, InSAR, holographic interferometry) are subject to aliasing, i.e., jumps that occur whenever the phase shift equals $\pm\pi$. Our study aimed to address the issue of phase jumps. Among the traditional methods for phase unwrapping, one can mention: *i*) graph cuts [19], [20], which involve representing the wrapped phase data as a graph and determining the minimum cut that separates the known and unknown phase values; *ii*) least-squares approaches, which minimize the differences between partial derivatives of the wrapped phase and those of the unwrapped solution [8], [21]. Specifically for color

TABLE V: 9-fold cross-validation final dealiasing results of DL methods with their best configurations against non-DL DeAN method.

Methods	Cosim $\pm\sigma$	Accuracy $\pm\sigma$	Recall $\pm\sigma$	Precision $\pm\sigma$
DeAN (Q=10)	0.95 ± 0.03	0.81 ± 0.08	0.62 ± 0.16	0.53 ± 0.20
DeAN (Optimized Q)	0.98 ± 0.01	0.91 ± 0.04	0.83 ± 0.08	0.80 ± 0.13
PDNet	0.98 ± 0.02	0.94 ± 0.06	0.87 ± 0.12	0.83 ± 0.10
nnU-Net	0.99 ± 0.01	0.96 ± 0.03	0.91 ± 0.06	0.89 ± 0.06
BATFormer	0.98 ± 0.02	0.91 ± 0.07	0.81 ± 0.14	0.91 ± 0.07

Doppler echocardiography, Muth *et al.* developed DeAN, a dealiasing algorithm based on statistical region merging [11], which was used in this study for comparative purposes.

Recently, deep learning (DL) techniques have been used to improve traditional methods in phase unwrapping [12], [22]–[24]. Our goal was to obtain alias-free color Doppler echocardiography by applying DL to dealias clinical Doppler velocity fields. DL-based approaches have been introduced for 2D phase unwrapping in InSAR [24]. Unlike echocardiographic images, InSAR interferograms are subject to multiple wraps, making the networks proposed in our study not suited since we focused only on single aliasing. On the other hand, InSAR images are not subject to significant clutter, whereas clutter in Doppler echocardiography generates substantial noise near moving tissues, making 2D phase unwrapping challenging. As a result, non-DL approaches such as graph cuts or least-squares methods, which work well for InSAR interferograms, are not effective for echocardiographic Doppler fields. Although the DeAN technique largely solved the problem, it still fails in some situations, as shown in our study. Therefore, we turned to DL and conducted an in-depth analysis and comparison of three architectures, including PDNet, which utilizes an unfolding framework. In addition, we illustrated the potential benefits of incorporating Doppler power information, since low power generally indicates poor blood Doppler signal. To better balance the aliased and non-aliased input data during training, we resorted to data augmentation by generating synthetic aliasing.

A. Comparison of the DL Methods

Our study found that the three DL methods we tested (PDNet, nnU-Net, and BATFormer) outperformed the non-DL DeAN method for color Doppler dealiasing. Notably, we observed that nnU-Net had the best performance, suggesting that the 2D U-Net architecture used in nnU-Net may be particularly well-suited for this task due to its ability to effectively capture spatial features. For example, in a challenging case where the aliased and non-aliased regions had similar hues (Fig. 4, third column), nnU-Net was able to unwrap correctly

while other DL methods failed or were less successful. Similar structures corrected by an expert were part of the training dataset, which implies that nnU-Net probably learned the flow patterns and leveraged this knowledge to achieve successful outcomes. PDNet also performed well, requiring > 200 times fewer parameters than nnU-Net, highlighting the potential for simpler DL models to achieve competitive results in color Doppler processing. Further exploration of this type of unfolding approach, including more complex modeling of the forward operator, is needed.

Although the third input strategy (velocity-power multiplication) contained less information than the second (velocity-power concatenation), it performed slightly better in the difficult fold (last row of Table II). The multiplication strategy largely suppressed velocity discontinuities in noisy regions, making the training task easier. In contrast, the concatenation of Doppler power and velocity allowed the model to learn the best strategy for combining these two inputs, which could be beneficial in larger datasets.

While we did not observe significant performance gains with BATFormer, the addition of synthetic data improved the outcomes, indicating that BATFormer also has potential for color Doppler dealiasing, especially when more data is available. Our results demonstrate that DL methods can significantly improve upon traditional methods for color Doppler processing. Additionally, they underscore the importance of further investigating the performance of different DL architectures for this task and finding ways to effectively exploit the strengths of each architecture.

B. Limitations and Perspectives

Color Doppler aliasing in the left ventricle mainly occurs in the mitral jet during early and late filling, as well as during ejection into the ventricular outflow tract. As depicted in the figures, the aliasing in our study was single. However, in certain valvular diseases, such as mitral stenosis or aortic regurgitation, multiple aliasing can occur in the intraventricular cavity due to the high fluctuating velocities of the turbulent jet. The nature of multiple aliasing in this context differs from that observed in InSAR, requiring specific studies to assess the feasibility of removing aliasing in areas with significant local flow perturbations. Although this remains to be verified, it is likely that a similar strategy could also work with disturbed flows, provided that we have access to Doppler data with their alias-free references. Such ground truths could be obtained by supervised correction, as in this study, and by simulations [25].

Since we used a clinical ultrasound system with a color Doppler rate of 10 to 15 frames per second, our study did not exploit temporal information. In the context of high-frame-rate echocardiography [26], neural networks with enforced temporal consistency [6] or 3D U-Net could potentially improve dealiasing performance by leveraging temporal information. This approach would be especially relevant as high-frame-rate color Doppler is subject to more noise related to clutter signals.

C. Applications in Quantitative Color Doppler

Once corrected, color Doppler images can be used to visualize and quantify intracardiac blood flow. As mentioned in the

introduction, intraventricular vector flow mapping (iVFM) is an approach to obtain comprehensive flow information, from which hemodynamic parameters can be estimated. Using a color M-mode, it is also possible to estimate the pressure difference between the apex and the mitral base, which reflects the cardiac filling [27], [28]. However, prior dealiasing is required for this method [28]. To this end, the approaches outlined in this study could be used with color M-mode images. In a more ambitious perspective, it would be conceivable to develop neural networks that can directly infer velocity vector fields or relative pressure fields from color Doppler images, once properly trained. In this case, the dealiasing process would be intrinsically integrated into the network. The main difficulty would lie in obtaining paired input data that provide the reference values. Simulations combining flows and acoustics could provide a relevant avenue for this purpose [25], [29].

V. CONCLUSION

While traditional methods are effective for interferometric imaging, they are limited for color Doppler echocardiography due to high noise generated by clutter signals in moving tissues surrounding the blood flow. We have demonstrated that deep learning (DL) techniques can achieve alias-free color Doppler echocardiography. Our proposed DL methods outperformed the non-DL DeAN method, with nnU-Net achieving the best performance, followed by PDNet. In addition, the incorporation of power information and artificial aliasing augmentation improved the results. The application of DL techniques to color Doppler echocardiography is a promising approach that could enhance the clinical utility of this widely used imaging modality.

ACKNOWLEDGMENT

The authors gratefully acknowledge financial support of the MEGA doctoral school (ED 162), the French National Research Agency (ANR) through the “4D-iVFM” project (ANR-21-CE19-0034-01), the LABEX PRIMES (ANR-11-LABX-0063) of Université de Lyon, and the LABEX CELYA (ANR-10-LABX-0060) within the program “Investissements d’Avenir” operated by the French ANR. The authors thank GENCI-IDRIS for providing access to HPC resources (Grant 2022-[AD010313603]).

REFERENCES

- [1] W. A. Zoghbi, D. Adams, R. O. Bonow *et al.*, “Recommendations for noninvasive evaluation of native valvular regurgitation: a report from the American Society of Echocardiography developed in collaboration with the Society for Cardiovascular Magnetic Resonance,” *Journal of the American Society of Echocardiography*, vol. 30, no. 4, pp. 303–371, 2017.
- [2] K. Hanumansetty, A. M. Aspari, and D. J. Palamattam, “Ventricular septal defects in two dimensional imaging seen as a single defect on colour flow doppler imaging,” *Annals of Cardiac Anaesthesia*, vol. 25, no. 2, p. 240, 2022.
- [3] M. V. Revzin, A. Imanzadeh, C. Menias *et al.*, “Optimizing image quality when evaluating blood flow at doppler us: a tutorial,” *Radiographics*, vol. 39, no. 5, pp. 1501–1523, 2019.
- [4] F. Vixège, A. Berod, Y. Sun *et al.*, “Physics-constrained intraventricular vector flow mapping by color Doppler,” *Physics in Medicine & Biology*, vol. 66, no. 24, p. 245019, 2021.
- [5] F. Vixège, A. Berod, P.-Y. Courand *et al.*, “Full-volume three-component intraventricular vector flow mapping by triplane color Doppler,” *Physics in Medicine and Biology*, vol. 67, no. 9, 2022.
- [6] N. Painchaud, N. Duchateau, O. Bernard *et al.*, “Echocardiography segmentation with enforced temporal consistency,” *IEEE Transactions on Medical Imaging*, vol. 41, no. 10, pp. 2867–2878, 2022.
- [7] D. Posada, J. Porée, A. Pellissier *et al.*, “Staggered Multiple-PRF Ultrafast Color Doppler,” *IEEE Transactions on Medical Imaging*, vol. 35, 2016.
- [8] D. C. Ghiglia and L. A. Romero, “Robust two-dimensional weighted and unweighted phase unwrapping that uses fast transforms and iterative methods,” *Journal of the Optical Society of America A*, vol. 11, no. 1, p. 107, 1994.
- [9] A. P. Shanker and H. Zebker, “Edgelist phase unwrapping algorithm for time series InSAR analysis,” *Journal of the Optical Society of America A*, vol. 27, no. 3, pp. 605–612, 2010.
- [10] J. Martinez-Carranza, K. Falaggis, and T. Kozacki, “Fast and accurate phase-unwrapping algorithm based on the transport of intensity equation,” *Applied Optics*, vol. 56, no. 25, pp. 7079–7088, 2017.
- [11] S. Muth, S. Dort, I. A. Sebag *et al.*, “Unsupervised dealiasing and denoising of color-doppler data,” *Medical Image Analysis*, vol. 15, no. 4, pp. 577–588, 2011.
- [12] H. Nahas, J. S. Au, T. Ishii *et al.*, “A deep learning approach to resolve aliasing artifacts in ultrasound color flow imaging,” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 67, no. 12, pp. 2615–2628, 2020.
- [13] J. Adler and O. Öktem, “Learned primal-dual reconstruction,” *IEEE Transactions on Medical Imaging*, vol. 37, no. 6, pp. 1322–1332, 2018.
- [14] D. P. Kingma and J. Ba, “Adam: A Method for Stochastic Optimization,” in *Proc. ICLR*, 2015.
- [15] F. Isensee, P. F. Jaeger, S. A. A. Kohl *et al.*, “nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation,” *Nature Methods*, vol. 18, no. 2, pp. 203–211, 2021.
- [16] M. Antonelli, A. Reinke, S. Bakas *et al.*, “The Medical Segmentation Decathlon,” *Nature Communications*, vol. 13, no. 1, p. 4128, 2022.
- [17] X. Lin, L. Yu, K.-T. Cheng *et al.*, “Batformer: Towards boundary-aware lightweight transformer for efficient medical image segmentation,” *IEEE Journal of Biomedical and Health Informatics*, pp. 1–13, 2023.
- [18] F. Mehregan, F. Tournoux, S. Muth *et al.*, “Doppler vortography: a color Doppler approach for quantification of the intraventricular blood flow vortices,” *Ultrasound in medicine & biology*, vol. 40, no. 1, pp. 210–221, 2014.
- [19] J. Dong, F. Chen, D. Zhou *et al.*, “Phase unwrapping with graph cuts optimization and dual decomposition acceleration for 3d high-resolution mri data,” *Magnetic resonance in medicine*, vol. 77, no. 3, pp. 1353–1358, 2017.
- [20] J. M. Bioucas-Dias and G. Valadao, “Phase unwrapping via graph cuts,” *IEEE Transactions on Image processing*, vol. 16, no. 3, pp. 698–709, 2007.
- [21] B. Hunt, “Matrix formulation of the reconstruction of phase values from phase differences,” *JOSA*, vol. 69, no. 3, pp. 393–399, 1979.
- [22] G. Spoorthi, S. Gorthi, and R. K. S. S. Gorthi, “Phasenet: A deep convolutional neural network for two-dimensional phase unwrapping,” *IEEE Signal Processing Letters*, vol. 26, no. 1, pp. 54–58, 2018.
- [23] K. Wang, Y. Li, Q. Kemao *et al.*, “One-step robust deep learning phase unwrapping,” *Optics express*, vol. 27, no. 10, pp. 15 100–15 115, 2019.
- [24] H. Zhou, C. Cheng, H. Peng *et al.*, “The phu-net: A robust phase unwrapping method for mri based on deep learning,” *Magnetic Resonance in Medicine*, vol. 86, no. 6, pp. 3321–3333, 2021.
- [25] Y. Sun, F. Vixège, K. Faraz *et al.*, “A pipeline for the generation of synthetic cardiac color doppler,” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 69, no. 3, pp. 932–941, 2022.
- [26] J. Faurie, M. Baudet, K. C. Assi *et al.*, “Intracardiac vortex dynamics by high-frame-rate doppler vortography—in vivo comparison with vector flow mapping and 4-d flow mri,” *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, vol. 64, no. 2, pp. 424–432, 2016.
- [27] K. Sørensen, S. Fadnes, L. Mertens *et al.*, “Assessment of early diastolic intraventricular pressure difference in children by blood speckle-tracking echocardiography,” *Journal of the American Society of Echocardiography*, 2023.
- [28] A. Hodzic, D. Garcia, E. Saloux *et al.*, “Echocardiographic evidence of left ventricular untwisting-filling interplay,” *Cardiovascular Ultrasound*, vol. 18, no. 1, pp. 1–12, 2020.
- [29] S. Shahriari and D. Garcia, “Meshfree simulations of ultrasound vector flow imaging using smoothed particle hydrodynamics,” *Physics in Medicine & Biology*, vol. 63, no. 20, p. 205011, 2018.