Neural Multi-Scale Self-Supervised Registration for Echocardiogram Dense Tracking

Wentao Zhu¹, Yufang Huang², Mani A Vannan³, Shizhen Liu³, Daguang Xu¹,
Wei Fan⁴, Zhen Qian⁴, and Xiaohui Xie⁵

¹ NVIDIA wentaoz1@ics.uci.edu

² Lenovo Research

³ Piedmont Heart Institute

⁴ Tencent Medical AI Lab

⁵ University of California, Irvine xhx@ics.uci.edu

Abstract. Echocardiography has become routinely used in the diagnosis of cardiomyopathy and abnormal cardiac blood flow. However, manually measuring myocardial motion and cardiac blood flow from echocardiogram is time-consuming and error-prone. Computer algorithms that can automatically track and quantify myocardial motion and cardiac blood flow are highly sought after, but have not been very successful due to noise and high variability of echocardiography. In this work, we propose a neural multi-scale self-supervised registration (NMSR) method for automated myocardial and cardiac blood flow dense tracking. NMSR incorporates two novel components: 1) utilizing a deep neural net to parameterize the velocity field between two image frames, and 2) optimizing the parameters of the neural net in a sequential multi-scale fashion to account for large variations within the velocity field. Experiments demonstrate that NMSR yields significantly better registration accuracy than the state-of-the-art methods, such as advanced normalization tools (ANTs) and VoxelMorph, for both myocardial and cardiac blood flow dense tracking. Our approach promises to provide a fully automated method for fast and accurate analyses of echocardiograms.

Keywords: Neural multi-scale self-supervised registration · Echocardiogram registration · Myocardial tracking · Cardiac blood flow tracking.

1 Introduction

Cardiovascular diseases are the leading cause of mortality and morbidity globally. Echocardiogram is a non-invasive ultrasonic imaging test that exams the heart in motion. Because of the relatively low cost and ease of access, it has become a routinely used diagnostic tool in cardiology. Echocardiogram provides a wealth of information. Besides the assessment of the geometry and pumping capacity of the heart, echocardiogram has been used for the evaluation of the global and regional myocardial contractility using speckle tracking techniques [8].

The use of contrast agents in echocardiogram has been proposed to further improve the imaging quality of echocardiogram. The advantages of using contrast include 1) a sharper demarcation between the left ventricular cavity and

the myocardium, and 2) easier detection of sludges and thrombi in the heart. In addition, a high-power ultrasound beam produces a swirling pattern in the contrast-enhanced blood pool, which facilitates the visualization of the ventricular flow vortex. Unlike the conventional color Doppler technique, which estimates the flow velocity along a specific angle, the vortex imaging in contrast echocardiogram enables flow visualization and quantification in 2D, which is an important hemodynamic indicator of the heart’s pumping performance and efficiency [1]. In this work, we aim to achieve dense tracking of the myocardium in non-contrast echocardiogram and the blood flow in contrast echocardiogram, in an effort to make echocardiogram a one-stop shop for both myocardial contractility assessment and ventricular flow pattern analysis.

There are many previous methods on deformable registration that can be utilized for dense tracking in echocardiograms [9]. Traditional registration methods are based on the optimization in registration field space such as elastic-type models [5, 18], FFD [17], Demons [20] and statistical parametric mapping [2]. Diffeomorphic transformations preserve topology and many methods are derived from them such as LDDMM [7] and SyN [4]. The optimization of these traditional methods typically require substantial time. Deep learning based registration methods usually rely on ground truth of registration field [15, 19]. Recent unsupervised deep learning based registrations, such as VoxelMorph [6], are facilitated by the spatial transformer network [21, 12], and the VoxelMorph is further extended to diffeomorphic transformation and Bayesian framework [10]. Adversarial similarity network adds an extra discriminator and uses adversarial training to improve the unsupervised registration [11]. These purely learning based methods cannot be directly applied to ultrasound images for velocity estimation especially for vortex detection in cardiac blood flow because of great noise in echocardiogram, large velocity variations of cardiac blood flow and large amounts of missing and new blood within the ultrasound plane.

Inspired by the improvement of multi-scale registration and neural network parameterized optimization [9], we propose a neural self-supervised optimization based multi-scale framework for dense tracking in echocardiograms as illustrated in Fig. 1. The neural self-supervised optimization yields accurate velocity field estimation by eliminating the gap between the estimation on the training set and that on the test set and alleviating the optimization difficulties such as local minima. Multi-scale strategy provides a sequential optimization pathway to the flow tracking that naturally emulates the formation of the fractal patterns of turbulent flow.

2 Neural Multi-Scale Self-Supervised Registration

We denote an ultrasound sequence by $\mathbf{I} = \{\mathbf{I}_1, \mathbf{I}_2, \dots, \mathbf{I}_t, \dots, \mathbf{I}_n\}$, where $\mathbf{I}_t \in \mathbb{R}^{h \times w \times c}$ is the t_{th} frame of $\mathbf{I}$, and h , w and c are image height, width and the number of channels respectively. In this work, we focus on calculating the registration field $\mathbf{F}_t$ as velocity field estimation which can be used by dense tracking for the two neighbor frames $\mathbf{I}_t$ and $\mathbf{I}_{t+1}$. We employ U-Net with skip

Fig. 1. Framework of the proposed neural multi-scale self-supervised registration. We employ U-Net to parameterize the registration function $\tilde{F}_t^s$ for velocity estimation between the reconstructed frame M_{t+1}^s from last scale $s/2$ and $t + 1_{th}$ frame I_{t+1}^s in the scale s . Neural self-supervised optimization with reconstruction loss and smoothness loss is conducted to yield robust and accurate velocity field for each scale.

connections to obtain the velocity field F_t [16]. The framework of neural multi-scale self-supervised registration (NMSR) is illustrated in Fig. 1.

Because the velocity of cardiac blood varies greatly, we extend the neural self-supervised registration to a multi-scale framework. We construct multi-scale U-Nets with scales $\{S, 2S, \dots, s, \dots, 1/2, 1\}$ to parameterize the registration fields of these scales where S is the coarsest scale. In the registration of each scale s , we try to obtain the registration field $\tilde{F}_t^s$ by which we transform the reconstructed frame $\mathcal{T}(I_t, F_t^{s/2})$ from the last scale $s/2$ to the $t + 1_{th}$ frame I_{t+1} , where $t = 1, \dots, n - 1$ and $\mathcal{T}$ is a spatial transformer network [12]. If s is the coarsest scale S , $\mathcal{T}(I_t, F_t^{s/2}) = I_t$. Given an ultrasound sequence I , we firstly resize the moving image $\mathcal{T}(I_t, F_t^{s/2})$ and the fixed image I_{t+1} to the current scale s for data preparation as

$$M_{t+1}^s = \mathcal{P}_1(\mathcal{T}(I_t, F_t^{s/2}), s), \quad I_{t+1}^s = \mathcal{P}_1(I_{t+1}, s), \quad (1)$$

where M_{t+1}^s is the scaled reconstructed frame and $\mathcal{P}_1$ is the down-sampling.

In addition to the high variations of velocities for cardiac blood flow, the signal-to-noise ratio of ultrasound image is typically low which easily leads to inaccurate prediction of the registration field. This aggregates the difficulty of learning and inference in the multi-scale registration model because the error can be accumulated over scales and the noisy error along the high dimensional output space makes the learning even harder than learning in the ultrasound space which is empirically observed in our experiments. We propose neural self-supervised optimization based on one ultrasound sequence, considering that the

optimization based methods can eliminate the undesirable generalization ability of purely learning based model, and the settings and motions are consistent within one ultrasound sequence. We use stochastic optimization for the neural parameterized registration to alleviate the optimization difficulties such as local minima in traditional registration methods [14]. The neural self-supervised optimization for the current scale s and ultrasound sequence $\mathbf{I}$ can be formulated

$$\begin{aligned}\tilde{\mathbf{F}}_t^s &= \mathbf{N}(\mathbf{M}_{t+1}^s, \mathbf{I}_{t+1}^s; \boldsymbol{\theta}^s), \\ \boldsymbol{\theta}_*^s &= \arg \min_{\boldsymbol{\theta}^s} \sum_{t=1}^{n-1} \mathcal{L}_R(\mathcal{T}(\mathbf{M}_{t+1}^s, \tilde{\mathbf{F}}_t^s), \mathbf{I}_{t+1}^s; \boldsymbol{\theta}^s) + \lambda \mathcal{L}_S(\tilde{\mathbf{F}}_t^s; \boldsymbol{\theta}^s),\end{aligned}\quad (2)$$

where $\mathbf{N}$ is the neural network with parameters $\boldsymbol{\theta}^s$, $\mathcal{L}_R$ is the reconstruction loss to measure the similarity between the moving image and fixed image, $\mathcal{L}_S$ is the smoothness loss for the registration field, λ is the trade-off between reconstruction loss and smoothness loss.

We use negative local cross-correlation loss as reconstruction loss and L_2 norm of registration field gradient as the smoothness loss. For clarity, we omit s in the reconstruction loss $\mathcal{L}_R$ and write the two losses as

$$\begin{aligned}\mathcal{L}_R(\mathbf{M}_t, \mathbf{I}_t; \boldsymbol{\theta}^s) &= - \sum_{\mathbf{p}} \frac{(\sum_{\mathbf{p}_i} (\mathbf{M}_t(\mathbf{p}_i) - \mathbf{M}'_t(\mathbf{p}))(\mathbf{I}_t(\mathbf{p}_i) - \mathbf{I}'_t(\mathbf{p})))^2}{\sum_{\mathbf{p}_i} (\mathbf{M}_t(\mathbf{p}_i) - \mathbf{M}'_t(\mathbf{p}))^2 \sum_{\mathbf{p}_i} (\mathbf{I}_t(\mathbf{p}_i) - \mathbf{I}'_t(\mathbf{p}))^2}, \\ \mathcal{L}_S(\tilde{\mathbf{F}}_t^s; \boldsymbol{\theta}^s) &= \|\nabla_x \tilde{\mathbf{F}}_t^s\|_{L_2} + \|\nabla_y \tilde{\mathbf{F}}_t^s\|_{L_2},\end{aligned}\quad (3)$$

where $\mathbf{p}$ is the pixel position in the frame and $\mathbf{p}_i$ is the pixel position within a square with the center as $\mathbf{p}$, $\mathbf{M}'_t(\mathbf{p})$ and $\mathbf{I}'_t(\mathbf{p})$ are local means of pixel position $\mathbf{p}_i$ in $\mathbf{M}_t$ and $\mathbf{I}_t$ respectively.

After the neural self-supervised optimization, we obtain the optimal parameters $\boldsymbol{\theta}_*^s$ of the neural network and can calculate the velocity field $\tilde{\mathbf{F}}_t^s$ for the reconstructed frame $\mathcal{T}(\mathbf{I}_t, \mathbf{F}_t^{s/2})$ from the last scale $s/2$. We calculate the registration field $\mathbf{F}_t^s$ for the ultrasound frame $\mathbf{I}_t$ by combining registration field $\mathbf{F}_t^{s/2}$ of the last scale and intermediate field $\tilde{\mathbf{F}}_t^s$. For each pixel position $\mathbf{p}$ in the ultrasound frame $\mathbf{I}_t$, we can obtain the final position $\hat{\mathbf{p}}$ and the combined registration field $\mathbf{F}_t^s$ by

$$\mathbf{F}_t^s(\mathbf{p}) = \hat{\mathbf{p}} - \mathbf{p} = \mathbf{F}_t^{s/2}(\mathbf{p}) + \mathcal{P}_2\left(\frac{1}{s}\tilde{\mathbf{F}}_t^s, \frac{1}{s}\right)(\mathbf{p} + \mathbf{F}_t^{s/2}(\mathbf{p})), \quad (4)$$

where $\mathcal{P}_2$ is the linear interpolation and we use linear interpolation to calculate the field for $\mathbf{p} + \mathbf{F}_t^{s/2}(\mathbf{p})$.

3 Experiments

We collected echocardiograms from 19 patients with 3,052 frames in total for myocardial tracking, and contrast echocardiograms from 71 patients with 11,462

Table 1. Comparisons on myocardial and cardiac blood dense tracking among ANTs, VoxelMorph and NMSR. NMSR obtains the best performance.

Methods	Myocardial tracking		Cardiac blood tracking	
	MSE (10^{-3})	Mean CC (10^{-1})	MSE (10^{-3})	Mean CC (10^{-1})
ANTs	15.5179±9.4637	3.1597±1.3913	3.9344±1.4660	4.2062±1.0576
VoxelMorph	1.2266±0.5457	4.7363±0.5457	5.8654±1.6943	3.3462±0.5891
NMSR	1.3438±0.6248	4.6270±0.5340	5.9511±1.6789	3.2228±0.5754
NMSR (1/8)	1.7938±0.5992	4.1428±0.5066	6.8603±1.4958	2.3314±0.5497
NMSR (1/4)	1.8185±0.5437	4.2746±0.4700	5.6694±1.3347	2.7566±0.6186
NMSR (1/2)	1.5405±0.4515	4.5508±0.4601	4.6870±1.0579	3.2854±0.7191
NMSR (1)	1.3204±0.4058	4.8954±0.4173	4.0337±0.9903	3.9781±0.6917
NMSR _V	1.2158±0.5473	4.7809±0.4952	5.7780±1.6222	3.3942±0.5868
NMSR _V (1/8)	1.6937±0.5616	4.2084±0.5086	6.8108±1.8189	2.3948±0.5495
NMSR _V (1/4)	1.7022±0.4434	4.3377±0.4674	5.6577±1.3317	2.7109±0.5988
NMSR _V (1/2)	1.4363±0.3625	4.6459±0.4451	4.6228±1.1381	3.3155±0.6691
NMSR _V (1)	1.2420±0.3041	5.0280±0.3922	3.9287±1.0209	4.1209±0.6188
NMSR _{VI} (1/4)	1.3504±0.4885	4.4140±0.5005	5.6106±1.2424	2.7331±0.6779
NMSR _{VI} (1/2)	1.0929±0.3775	4.7155±0.4569	4.5089±1.0272	3.3879±0.7524
NMSR _{VI} (1)	0.9206±0.3236	5.0881±0.4107	3.7944±0.9384	4.2519±0.7181

frames in total for cardiac blood tracking. For testing, we randomly choose three patients' echocardiograms with 291 frames for myocardial tracking, and three patients' echocardiograms with 216 frames for cardiac blood tracking from the two datasets. The rest of echocardiograms are used as training. All echocardiograms have no registration field ground truth.

We only use the first channel of echocardiography images (e.g., treated as gray-scale images) with the pixel values normalized to be in $[0, 1]$. For cardiac blood tracking, we extract cardiac blood region by 1) creating masks of the left ventricular blood pool at the end of the systole and the end of the diastole, 2) using active contour model to fit 100 uniformly sampled spline points along a circle into the boundary of cardiac blood mask [13], 3) using linear interpolation to get 100 interpolated spline points for each frame, 4) using radial basis function in interpolation to get the final smooth cardiac blood boundary from the 100 spline points. Removing myocardial region is crucial to cardiac blood tracking.

For unsupervised learning or self-supervised optimization, one of the main challenges is the model evaluation. Manually labeling the corresponding points for evaluation is time-consuming, laborious and inaccurate, because the total size of one frame is 1024×768 and the signal-to-noise ratio is low. Instead of using pixel position based evaluation metric, we use reconstruction based metrics, i.e., the mean square error (MSE) and the mean local cross correlation (Mean CC) with radius as 10. For the metrics in Table 1, we calculate the average MSE and Mean CC over all frame pairs I_t and I_{t+1} with the pixel value of range $[0, 1]$. For MSE and Mean CC of one frame pair, we take the average of square error and local cross correlation over the masked region.

We compare our approach to symmetric normalization (SyN) with cross-correlation similarity as implemented in ANTs [4, 3], and VoxelMorph [6]. ANTs (SyN) is a traditional optimization based method and VoxelMorph is a deep learning based unsupervised registration with a similar network structure and loss function as NMSR. For the purpose of ablation studies, we report the results of **1)** NMSR, using only one scale and neural self-supervised optimization, **2)** NMSR (1/8), the coarsest scale of neural multi-scale self-supervised registration, **3)** NMSR (1/4), the second coarsest scale of NMSR, **4)** NMSR (1/2), the third coarsest scale of NMSR, **5)** NMSR (1), the final scale of NMSR after sequential multi-scale optimization, **6)** NMSR_V, NMSR_V (*), which use training-data optimized NMSR as an initialization and conduct self-supervised neural optimization afterwards, and **7)** NMSR_{VI} (*), which use the neural weights from the last scale as an initialization.

For all these multi-scale based methods, we use four different scales 1/8, 1/4, 1/2 and 1. We use radius of 6 pixels for the local cross-correlation loss in all these methods. We set the number of optimization steps to 200 for each scale in ANTs. We use learning rate of 1×10^{-3} and Adam optimizer to update the weights in neural networks for both NMSR and VoxelMorph [14]. The λ is set to be 10. We set the number of optimization steps to 3500 per ultrasound sequence for the self-supervised neural optimization, and set the number of iterations to $3500 \times$ the number of training ultrasound sequences for VoxelMorph.

Quantitative comparison results of our model against the state-of-the-arts on both myocardial and cardiac blood flow dense tracking are shown in Table 1. We highlight the following a few observations: 1) NMSR_{VI} (1) achieves the best performances, and outperforms ANTs in terms of both MSE and mean CC on both tasks, likely due to the representation and optimization efficiency of deep neural nets; 2) NMSR_V yields consistently better results than VoxelMorph, demonstrating the efficacy of self-supervised optimization during the test phase for improving velocity field estimation and reducing the estimation gap between training and testing; 3) NMSR (1) achieves better performance than NMSR on all experiments, demonstrating the benefit of sequential multi-scale optimization in echocardiogram registration. The multi-scale scheme alleviates the over-optimization of reconstruction loss, which can be visually noticed from Fig. 2 and 3; and 4) NMSR_V (1) obtains better performance than NMSR (1), illustrating the benefit of using pretrained models as an initialization for NMSR, further confirmed by the better performance of NMSR_{VI} over NMSR_V.

Visualizations We visualize the myocardial tracking results based on ANTs, VoxelMorph and the intermediate registration fields from NMSR_{VI} with four different scales in Fig. 2. We randomly choose one frame from these ultrasound images. From Fig. 2, we note that the registration field from ANTs is noisy, and the velocity direction from VoxelMorph for the right myocardial is incorrect. By contrast, NMSR (1/8) produces the smoothest registration field, and NMSR (1) generates more detailed velocity estimation that preserves both large and low-scale velocity variations. The coarse-to-fine results illustrate that the multi-scale

Fig. 2. Visualization of myocardial tracking based on ANTs, VoxelMorph and NMSR. The right side of velocity field by VoxelMorph is inaccurate. NMSR yields smooth, coarse-to-fine and accurate velocity estimation.

optimization scheme coupled with deep neural nets can be very effective in dealing with the highly challenging case of image registration in echocardiograms.

We visualize the cardiac blood tracking results in Fig. 3. From the registration fields from ANTs and VoxelMorph, we cannot easily recognize the vortex in cardiac blood flow. By contrast, the vortex flow pattern from NMSR is readily recognizable. The general vortex pattern is apparent from the coarsest level registration by NMSR (1/8), followed by finer-scale registrations to introduce details of local velocity field variations. The final velocity field produced by NMSR (1) includes both easily recognizable vortex flow, as well as details of local field variations.

Computational Cost For ANTs, the average computational time is 214.10 ± 54.04 seconds for the registration of two consecutive frames on 12 processors of Intel i7-6850K CPU @ 3.60GHz. For an ultrasound sequence of 50 frames, the computational time is about **three hours** for ANTs. Because the inference of VoxelMorph only relies on one feed forward pass of deep neural network, the average computational time is 0.11 ± 0.47 seconds for one pair frames on one NVIDIA 1080 Ti GPU. The NMSR takes 279.97, 101.65, 68.79, 66.09 seconds for neural self-supervised optimization with the scale 1, 1/2, 1/4, 1/8 respectively on one ultrasound sequence of 49 frames by one NVIDIA 1080 Ti GPU. The

Fig. 3. Visualization of cardiac blood tracking based on ANTs, VoxelMorph and NMSR. ANTs and VoxelMorph produce noisy velocity fields. NMSR successfully detects the vortex and yields coarse-to-fine and smooth velocity field.

NMSR takes less than **nine minutes** on the neural self-supervised optimization in total for one ultrasound sequence, achieving **20** times speedup over ANTs.

4 Conclusion

In this work, we propose a novel framework, neural multi-scale self-supervised registration (NMSR), for both myocardial and cardiac blood dense tracking. To produce accurate velocity estimation from noisy ultrasound images and reduce the estimation gap between training and testing, we incorporate self-supervised optimization in the registration framework. To handle large variations of velocity fields in echocardiogram tracking, a multi-scale scheme is integrated into the proposed framework to reduce the over-optimization of similarity functions. Our proposed method consistently outperforms state-of-the-art methods on both myocardial and cardiac blood flow dense tracking. With further improvements on model and optimization, to consider for example other loss functions and extend it to diffeomorphic registrations, it seems plausible to have a fully automated method for echocardiogram analysis.

References

1. Abe, H., et al.: Contrast echocardiography for assessing left ventricular vortex strength in heart failure: a prospective cohort study. *EHJCI* **14**(11) (2013)
2. Ashburner, J., et al.: Voxel-based morphometry—the methods. *Neuroimage* (2000)
3. Avants, B.B., Tustison, N., Song, G.: Advanced normalization tools (ants) (2009)
4. Avants, B.B., et al.: Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *MIA* **12**(1), 26–41 (2008)
5. Bajcsy, R., Kovačič, S.: Multiresolution elastic matching. *CVGIP* (1989)
6. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: An unsupervised learning model for deformable medical image registration. In: *CVPR* (2018)
7. Beg, M.F., Miller, M.I., Trounev, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *IJCV* **61**(2), 139–157 (2005)
8. Carasso, S., et al.: Velocity vector imaging: standard tissue-tracking results acquired in normals—the vvi-strain study. *JASE* **25**(5), 543–552 (2012)
9. Curiale, A.H., et al.: Influence of ultrasound speckle tracking strategies for motion and strain estimation. *Medical image analysis* **32**, 184–200 (2016)
10. Dalca, A.V., Balakrishnan, G., Guttag, J., Sabuncu, M.R.: Unsupervised learning for fast probabilistic diffeomorphic registration. In: *MICCAI* (2018)
11. Fan, J., Cao, X., Xue, Z., Yap, P.T., Shen, D.: Adversarial similarity network for evaluating image alignment in deep learning based registration. In: *MICCAI* (2018)
12. Jaderberg, M., et al.: Spatial transformer networks. In: *NIPS*. pp. 2017–2025 (2015)
13. Kass, M., Witkin, A., Terzopoulos, D.: Snakes: Active contour models. *IJCV* (1988)
14. Kingma, D.P., et al.: Adam: A method for stochastic optimization. In: *ICLR* (2015)
15. Rohé, M.M., Datar, M., Heimann, T., Sermesant, M., Pennec, X.: Svf-net: learning deformable image registration using shape matching. In: *MICCAI* (2017)
16. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *MICCAI*. pp. 234–241. Springer (2015)
17. Rueckert, D., et al.: Nonrigid registration using free-form deformations: application to breast mr images. *IEEE transactions on medical imaging* **18**(8), 712–721 (1999)
18. Shen, D., Davatzikos, C.: Hammer: hierarchical attribute matching mechanism for elastic registration. *IEEE transactions on medical imaging* **21**(11), 1421 (2002)
19. Sokooti, H., et al.: Nonrigid image registration using multi-scale 3d convolutional neural networks. In: *MICCAI*. pp. 232–239. Springer (2017)
20. Thirion, J.P.: Image matching as a diffusion process: an analogy with maxwell’s demons. *Medical image analysis* **2**(3), 243–260 (1998)
21. de Vos, B.D., et al.: End-to-end unsupervised deformable image registration with a convolutional neural network. In: *DLMIA/LCDS*, pp. 204–212. Springer (2017)