



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.

WiNet: Wavelet-based Incremental Learning for Efficient Medical Image Registration

Xinxing Cheng¹, Xi Jia¹, Wenqi Lu², Qiufu Li³, Linlin Shen³,
Alexander Krull¹, and Jinming Duan¹(✉)

¹ School of Computer Science, University of Birmingham, Birmingham, B15 2TT, UK
j.duan@cs.bham.ac.uk

² Department of Computing and Mathematics,
Manchester Metropolitan University, Manchester, M15 6BH, UK

³ National Engineering Laboratory for Big Data System Computing Technology,
Shenzhen University, 518060, China

Abstract. Deep image registration has demonstrated exceptional accuracy and fast inference. Recent advances have adopted either multiple cascades or pyramid architectures to estimate dense deformation fields in a coarse-to-fine manner. However, due to the cascaded nature and repeated composition/warping operations on feature maps, these methods negatively increase memory usage during training and testing. Moreover, such approaches lack explicit constraints on the learning process of small deformations at different scales, thus lacking explainability. In this study, we introduce a model-driven WiNet that incrementally estimates scale-wise wavelet coefficients for the displacement/velocity field across various scales, utilizing the wavelet coefficients derived from the original input image pair. By exploiting the properties of the wavelet transform, these estimated coefficients facilitate the seamless reconstruction of a full-resolution displacement/velocity field via our devised inverse discrete wavelet transform (IDWT) layer. This approach avoids the complexities of cascading networks or composition operations, making our WiNet an explainable and efficient competitor with other coarse-to-fine methods. Extensive experimental results from two 3D datasets show that our WiNet is accurate and GPU efficient. Code is available at <https://github.com/x-xc/WiNet>.

Keywords: Efficient Deformable Image Registration · Diffeomorphic · (Inverse) Discrete Wavelet Transform · Incremental Learning

1 Introduction

Deformable image registration is an essential step in many medical imaging applications [19,24]. Given a moving and fixed image pair, deformable registration estimates a dense non-linear deformation field that aligns the corresponding anatomical structures. Conventional methods such as LDDMM [5], DARTEL [1], SyN [2], Demons [22] and ADMM [21] are time-consuming and computationally expensive, due to instance-level (pair-wise) iterative optimization. Nevertheless,

such methods may involve sophisticated hyperparameter tuning, limiting their applications in large-scale volumetric registration.

U-Net-based methods have recently dominated medical image registration due to their fast inference speed. Following the generalized framework of Voxel-Morph [4], some methods have been proposed by adding stronger constraints over the deformation field such as inverse/cycle consistency [25,16] and diffeomorphisms [9,8]. Another track of methods proposed more advanced neural blocks such as vision-transformer [7,6] to model long-range information but inevitably increases the computational cost, sacrificing the training and testing efficiency. To reduce the repeated convolution operations in U-Net architecture and fasten the speed, model-driven methods such as B-Spline [18] and Fourier-Net [13] have proposed to learn a low-dimensional representation of the displacement. However, displacement fields are not necessarily smooth. Therefore, learning low-frequency components of the displacement alone [13] or interpolating the displacement from a set of regular control points [18] may not be able to fully and accurately reconstruct the entire displacement field, limiting its applications to large and complex registrations. Additionally, discrete Fourier Transform (DFT) gives a global full-resolution analysis in the frequency domain, therefore, it is straightforward to extract low-frequency signals using central patches/crops as in [13]. However, learning patch-based high frequencies with DFT is difficult via [13] because, unlike low frequencies that can be shifted in the middle of the frequency spectrum, high frequencies are not confined to a rectangular area.

Recent works [10,27,26,17,23,12,15] progressively estimate the large and complex deformations using either multiple cascades (where each cascade estimates a small decomposition of the final deformation) or pyramid **coarse-to-fine** compositions of multi-scale displacements within one network. Though the cascaded methods [10,27,26] show improved registration accuracy for estimating large deformation by composing small deformations, their computational costs increase exponentially with more cascades. Pyramid methods [17,15,12,23], on the other hand, estimate multi-scale deformations and sequentially compose them to the final deformation. For instance, LapIRN [17] utilizes a Laplacian pyramid network to capture large deformations by composing three different scale flows. However, the optimal performance for LapIRN requires sophisticated iterative training of different scales, negatively affecting its training efficiency. PRNet++ [15] employs a dual-stream pyramid network for coarse-to-fine registration through sequential warping on multi-scale feature maps, while the adaptation of local 3D correlation layers massively increases memory usage and computational costs. ModeT [23] has introduced a motion decomposition transformer that utilizes neighborhood attention mechanisms to first estimate multi-head multi-scale deformations from two-stream hierarchical feature maps. It then employs weighting modules to fuse multi-head flows in each scale and generates the final deformation by composing the fused output from all scales.

In this work, we propose a model-driven multi-scale registration network by embedding the discrete Wavelet transform (DWT) as prior knowledge, which we term WiNet. Compared to model-driven B-Spline [18] and Fourier-Net [13],

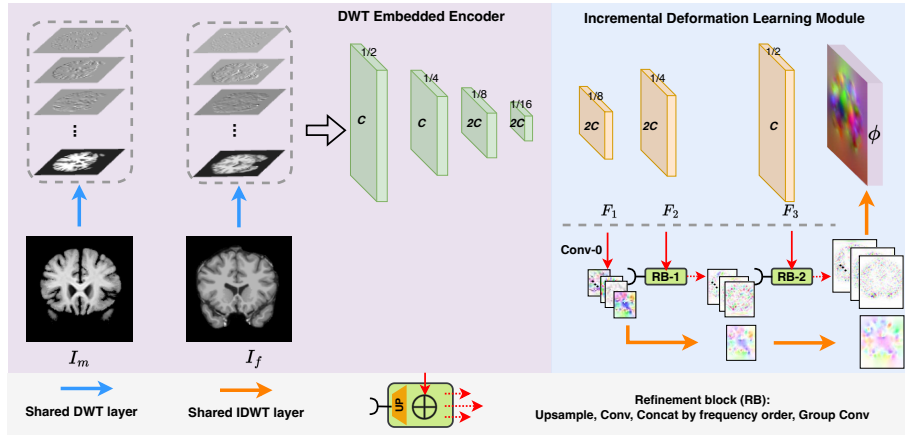


Fig. 1. Architecture of WiNet. Its encoder includes a shared DWT layer and four convolutional contracting layers. The decoder consists of three convolutional expansion layers and an Incremental Deformation Learning Module (which converts the convolutional features into final displacement ϕ with a conv-layer (Conv-0), two refinement blocks (RB-1 and RB-2), and a shared parameter-free IDWT layer). Note that the 3D images, features, and displacements are shown in 2D for illustration.

our WiNet learns a series of DWT coefficients that inherently preserve both low and high frequencies across different sub-bands, which have lower spatial resolution while maintaining overall structural information. Unlike pyramid methods [17,15,12,23], which rely solely on optimizing unsupervised training loss to estimate final deformation without explicitly constraining the learning process of small deformations at each scale, WiNet incrementally learns scale-wise DWT coefficients and naturally forms full-resolution deformation with a model-driven IDWT layer. Moreover, those pyramid methods require significant memory usage as they need to co-register multi-scale high-dimensional features, preventing their application on GPUs with less memory. In contrast, WiNet avoids compositions on feature maps and estimates only low-resolution DWT coefficients from low-resolution frequency components, making it memory-efficient.

WiNet contains a convolutional encoder-decoder network and an incremental deformation learning module designed to explicitly learn distinct frequency coefficients at various scales. Specifically, the initial decoding layer simultaneously captures both low-frequency and high-frequency coefficients at the lowest scale. The remaining decoding layers and the incremental module jointly refine the high-frequency coefficients, leading to a novel and explainable registration framework. We summarize our contributions as:

- We embed a differentiable DWT layer before the convolutional encoder, which empowers our network to operate on the low-resolution representation of images in various frequency bands.

- We introduce an incremental module that features three levels of IDWT for coarse-to-fine deformation estimation. This process takes advantage of both multi-scale prediction and DWT properties, avoiding the conventional multiple compositions and enhancing flexibility and interpretability.
- Extensive results on 3D brain and cardiac registration show that our WiNet can achieve comparable accuracy with state-of-the-art pyramid methods such as LapIRN [17], PRNet++ [15], and ModeT [23] while using only 31.9%, 25.6%, and 23.2% of their memory footprints (see Fig. 3), respectively.

2 Methodology

As illustrated in Figure 1, WiNet takes a moving I_m and fixed I_f pair to predict the full-resolution deformation ϕ . The network comprises two parts: 1) in the encoder, the DWT layer first decomposes each 3D input image into *eight* components at various frequency sub-bands using the defined orthogonal wavelet filter, then the convolutional layers learn hierarchical feature representations from the decomposed components; 2) in the decoder, the convolutional layers learn deformation related features at three scales. The incremental learning module takes the learned features from each decoding layer, progressively estimates scale-wise wavelet coefficients, and composes them into ϕ .

2.1 DWT Embedded Encoder

Our encoder includes a parameter-free differentiable DWT layer and four convolutional contracting layers. The DWT layer decomposes an image $\mathbf{I} \in \mathbb{R}^{D \times H \times W}$ by applying the low-pass \mathbf{L} and high-pass \mathbf{H} filters of an orthogonal Wavelet transform along the H , W , and D dimensions, respectively. Therefore, \mathbf{I} can be decomposed into one low-frequency component (\mathbf{I}_{lll}) and seven high-frequency components ($\mathbf{I}_{llh}, \mathbf{I}_{lhl}, \mathbf{I}_{lhh}, \mathbf{I}_{hll}, \mathbf{I}_{hhl}, \mathbf{I}_{hhl}, \mathbf{I}_{hhh}$):

$$\begin{aligned}
 \mathbf{I}_{lll} &= \mathbf{L} \oplus \mathbf{L} \oplus \mathbf{L} \oplus \mathbf{I}, & \mathbf{I}_{llh} &= \mathbf{L} \oplus \mathbf{L} \oplus \mathbf{H} \oplus \mathbf{I}, & \mathbf{I}_{lhl} &= \mathbf{L} \oplus \mathbf{H} \oplus \mathbf{L} \oplus \mathbf{I}, \\
 \mathbf{I}_{lhh} &= \mathbf{L} \oplus \mathbf{H} \oplus \mathbf{H} \oplus \mathbf{I}, & \mathbf{I}_{hll} &= \mathbf{H} \oplus \mathbf{L} \oplus \mathbf{L} \oplus \mathbf{I}, & & \\
 \mathbf{I}_{hhl} &= \mathbf{H} \oplus \mathbf{L} \oplus \mathbf{H} \oplus \mathbf{I}, & \mathbf{I}_{hhl} &= \mathbf{H} \oplus \mathbf{H} \oplus \mathbf{L} \oplus \mathbf{I}, & \mathbf{I}_{hhh} &= \mathbf{H} \oplus \mathbf{H} \oplus \mathbf{H} \oplus \mathbf{I},
 \end{aligned} \tag{1}$$

where \otimes , \oplus and \ominus denote matrix multiplication in H , W , and D dimensions, respectively; \mathbf{L} and \mathbf{H} are finite filters of wavelets that have truncated half-sizes for the dimension under filtering, e.g., $\lfloor \frac{H}{2} \rfloor \times H$ for the H dimension.

$$\mathbf{L} = \begin{pmatrix} \dots & \dots & \dots & & & & \\ \dots & l_{-1} & l_0 & l_1 & \dots & & \\ & & \dots & l_{-1} & l_0 & l_1 & \dots \\ & & & \dots & \dots & & \dots \end{pmatrix}, \quad \mathbf{H} = \begin{pmatrix} \dots & \dots & \dots & & & & \\ \dots & h_{-1} & h_0 & h_1 & \dots & & \\ & & \dots & h_{-1} & h_0 & h_1 & \dots \\ & & & \dots & \dots & & \dots \end{pmatrix}, \tag{2}$$

the values of \mathbf{L} and \mathbf{H} may vary based on the different orthogonal wavelet transform [20], while in this work, we only experiment with Haar and Daubechies.

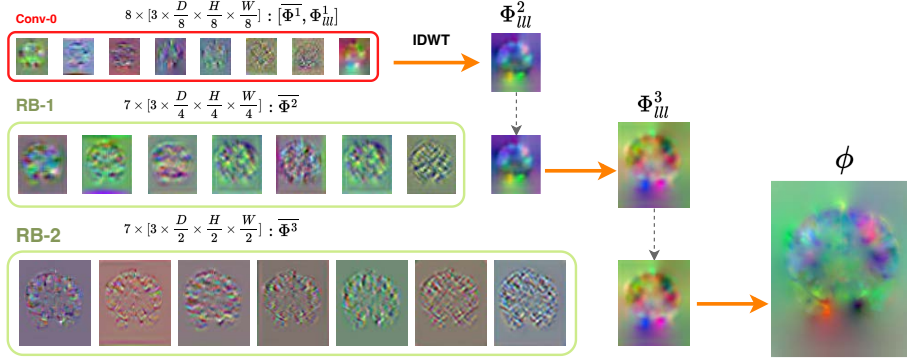


Fig. 2. Illustration of the coarse-to-fine incremental deformation learning module.

As shown the gray dashed box in Fig. 1, the DWT layer decomposes each of the full-resolution I_m and I_f into *eight* frequency components, we then concatenate them together to $[16, \frac{D}{2}, \frac{H}{2}, \frac{W}{2}]$ as the input of the following convolutional layers. The resolution of feature maps in each contracting convolution layer is $[C, \frac{D}{2}, \frac{H}{2}, \frac{W}{2}]$, $[C, \frac{D}{4}, \frac{H}{4}, \frac{W}{4}]$, $[2C, \frac{D}{8}, \frac{H}{8}, \frac{W}{8}]$, and $[2C, \frac{D}{16}, \frac{H}{16}, \frac{W}{16}]$, respectively.

The benefits of the embedded DWT layer in our encoder are twofold. First, the decomposed various frequency sub-bands potentially improve the network’s ability to capture both fine and coarse details. Second, it naturally halves the spatial resolution of the I_m and I_f without sacrificing information, saving computational costs for subsequent convolutional layers. We note that we made both the DWT and IDWT layers differentiable, while only the gradients of the IDWT layer are necessary in weight updating.

2.2 Incremental Deformation Learning Module & Decoder

The decoder consists of three expansive convolutional layers and an incremental deformation learning module. First, the convolutional decoding layers learn three different feature maps F_1, F_2 and F_3 with three different scales $[2C, \frac{D}{8}, \frac{H}{8}, \frac{W}{8}]$, $[2C, \frac{D}{4}, \frac{H}{4}, \frac{W}{4}]$ and $[C, \frac{D}{2}, \frac{H}{2}, \frac{W}{2}]$, respectively. The incremental learning module will learn for each scale *eight* different frequency sub-bands wavelet coefficients, which are denoted as $\Phi^i_{ll}, \Phi^i_{lh}, \Phi^i_{hl}, \Phi^i_{hh}, \Phi^i_{ll}, \Phi^i_{lh}, \Phi^i_{hl}, \Phi^i_{hh}$, where $i = 1, 2, 3$ represents the $\frac{1}{8}, \frac{1}{4}$ and $\frac{1}{2}$ scale, respectively. Specifically, $\Phi^{i=1}_{ll}$ is the low-frequency coefficient, while the rest seven coefficients, denoted as Φ^i , contain high-frequency details.

As depicted in Figs. 1 and 2, the Conv-0 in our incremental module initially estimates *eight* coefficients in the resolution of $[3, \frac{D}{8}, \frac{H}{8}, \frac{W}{8}]$ from feature F_1 to capture a coarse deformation. Then, the low-frequency Φ^2_{ll} is produced via the IDWT layer: $\Phi^2_{ll} = \text{IDWT}(\Phi^1) = \text{IDWT}([\Phi^1_{ll}, \Phi^1])$, where the IDWT operation,

according to Eqs. (1) and (2), can be denoted as

$$\begin{aligned} \Phi_{ll}^2 = & \mathbf{L} \oplus \mathbf{L} \oplus (\mathbf{L} \oplus \Phi_{ll}^1 + \mathbf{H} \oplus \Phi_{hl}^1) + \mathbf{L} \oplus \mathbf{H} \oplus (\mathbf{L} \oplus \Phi_{lh}^1 + \mathbf{H} \oplus \Phi_{hh}^1) \\ & + \mathbf{H} \oplus \mathbf{L} \oplus (\mathbf{L} \oplus \Phi_{hl}^1 + \mathbf{H} \oplus \Phi_{hh}^1) + \mathbf{H} \oplus \mathbf{H} \oplus (\mathbf{L} \oplus \Phi_{lh}^1 + \mathbf{H} \oplus \Phi_{hh}^1). \end{aligned} \quad (3)$$

Similarly, Φ_{ll}^3 can be produced by performing IDWT with Φ_{ll}^2 and the rest *seven* $[3, \frac{D}{4}, \frac{H}{4}, \frac{W}{4}]$ high-frequency coefficients, i.e., $\Phi_{ll}^3 = \text{IDWT}([\Phi_{ll}^2, \overline{\Phi^2}])$, then the final ϕ can be obtained as $\phi = \text{IDWT}([\Phi_{ll}^3, \overline{\Phi^3}])$. The primary challenge here is how can we learn the seven high-frequency $\overline{\Phi^2}$ and $\overline{\Phi^3}$ using $\overline{\Phi^1}$?

To address the challenge, we propose two refinement blocks (RB-1 and 2). As in Fig. 1, RB-1 incrementally refines $\overline{\Phi^1}$ to $\overline{\Phi^2}$ and RB-2 refines $\overline{\Phi^2}$ to $\overline{\Phi^3}$. Specifically, 1) the up-sampling operation in RB-1 (the yellow trapezoid) first converts the seven $[3, \frac{D}{8}, \frac{H}{8}, \frac{W}{8}]$ coefficients of $\overline{\Phi^1}$ into the resolution of $[3, \frac{D}{4}, \frac{H}{4}, \frac{W}{4}]$; 2) the conv layer (red arrow) takes learned feature F_2 and estimates also seven $[3, \frac{D}{4}, \frac{H}{4}, \frac{W}{4}]$ coefficients, which serve as the complementary residuals to the up-sampled coefficients; 3) then both the upsampled and learned coefficients are concatenated together in the frequency order; 4) and a group convolution (group=7, illustrated as the red dashed arrows) is used to output the refined $\overline{\Phi^2}$. The process can be denoted as

$$\overline{\Phi^2} = GC(Cat[Conv(F_2), UP(\overline{\Phi^1})]), \quad (4)$$

where UP , $Conv$, Cat , and GC denote upsampling, convolution, concatenation by frequency order, and group convolution, respectively. Likewise, we estimate $\overline{\Phi^3}$ from $\overline{\Phi^2}$ and F_3 with RB-2. Following the same procedure, the full-resolution deformation is produced by $\phi = \text{IDWT}([\Phi_{ll}^3, \overline{\Phi^3}])$.

2.3 Loss Functions

For diffeomorphic registration (WiNet-Diff), the output of WiNet-Diff is the stationary velocity fields \mathbf{v} and we employ seven scaling and squaring layers [1] to integrate the deformation field $\phi = \exp(\mathbf{v})$. The training loss of our network is $\mathcal{L}(\theta) = \mathcal{L}_{Sim}(I_m \circ (\phi(\theta) + \text{Id}), I_f) + \lambda \|\nabla \phi(\theta)\|_2^2$ or $\mathcal{L}(\theta) = \mathcal{L}_{Sim}(I_m \circ \exp(\mathbf{v}(\theta)), I_f) + \lambda \|\nabla \mathbf{v}(\theta)\|_2^2$, where θ represents the learnable network parameters, \circ refers the warping operator, Id is the identity grid, ∇ is the first order gradient and λ is the weight of the regularization term. The term $\mathcal{L}_{sim}(\cdot)$ quantifies the similarity between the warped moving image $I_m \circ \phi$ and the fixed image I_f . The second term is the smoothness regularization on ϕ .

3 Experiments

Datasets. One 3D brain MRI dataset (IXI) and one 3D cardiac MRI dataset (3D-CMR) were utilized. The pre-processed IXI⁴ dataset [6] comprises 576 MRI

⁴ <https://brain-development.org/ixi-dataset/>




Fig. 3. Comparison of the number of parameters (the area of circles), GPU training memory on IXI, and averaged GPU inference time on the same machine. Fourier-Net, TM-B-Spline-Diff, and B-Spline-Diff are abbreviated as F-Net, TM-BS-Diff, and BS-Diff, respectively. Our WiNet requires the lowest memory (6977 MiB) for training.

Table 1. Registration performance comparison on the IXI and 3D-CMR datasets.

Method	IXI		3D-CMR		
	Dice \uparrow	$ J _{<0\%}$	Dice \uparrow	HD \downarrow	$ J _{<0\%}$
Affine	.386 \pm .195	-	.493 \pm .043	8.40 \pm 0.89	-
SyN [3]	.645 \pm .152	<0.0001	.721 \pm .051	6.98 \pm 1.23	1.388 \pm 0.467
VM-1 [4]	.728 \pm .129	1.590 \pm 0.339	.723 \pm .032	6.53 \pm 1.00	0.783 \pm 0.325
VM-2 [4]	.732 \pm .123	1.522 \pm 0.336	.723 \pm .032	6.71 \pm 1.04	0.695 \pm 0.255
TM [6]	.754 \pm .124	1.579 \pm 0.328	.732 \pm .029	6.51 \pm 0.99	0.811 \pm 0.231
TM-Diff [6]	.594 \pm .163	<0.0001	.630 \pm .029	8.15 \pm 0.80	0.0\pm0.0
TM-B-Spline-Diff [6]	.761 \pm .122	<0.0001	.810 \pm .026	5.84\pm0.73	<0.0001
B-Spline-Diff [18]	.742 \pm .128	<0.0001	.809 \pm .033	6.04 \pm 0.73	0.0\pm0.0
Fourier-Net [13]	.763 \pm .129	0.024 \pm 0.019	.814 \pm .024	6.00 \pm 0.77	2.137 \pm 0.658
Fourier-Net-Diff [13]	.761 \pm .131	0.0\pm0.0	.827 \pm .027	6.08 \pm 0.85	<0.0003
LapIRN [17]	.760 \pm .127	0.312 \pm 0.106	.760 \pm .028	6.61 \pm 0.90	1.100 \pm 0.299
PRNet++ [15]	.755 \pm .130	1.052 \pm 0.302	.749 \pm .031	6.74 \pm 1.02	0.257 \pm 0.115
ModeT [23]	.758 \pm .125	0.114 \pm 0.057	.774 \pm .030	6.72 \pm 0.96	0.041 \pm 0.017
WiNet	.763 \pm .125	0.874 \pm 0.284	.814 \pm .027	6.06 \pm 0.66	2.160 \pm 0.612
WiNet-Diff	.765\pm.125	0.0\pm0.0	.832\pm.023	5.86 \pm 0.68	0.007 \pm 0.007

scans ($160 \times 192 \times 224$) from healthy subjects. We followed the same protocol with partitions 403, 58, and 115 for the training, validation, and testing sets, respectively. Atlas-based brain registration was conducted for each scan using a generated atlas from [16]. The 3D-CMR [21,14,11] contains 220 pairs of End-diastolic (ED) and End-systolic (ES) scans. All MRI scans were resampled to the same resolution $1.2 \times 1.2 \times 1.2 \text{ mm}^3$ and center-cropped to dimensions of $128 \times 128 \times 96$. We randomly split it into 100 training, 20 validation, and 100 testing sets.




Fig. 4. Visual comparisons on IXI and 3D-CMR. Column 1: Fixed image, Moving image. Columns 2-10: warped moving images, displacement fields as RGB images.

Evaluation Metrics. We use the Dice score, Hausdorff distance (HD), and the percentage of negative values of the Jacobian determinant ($|J|_{<0}\%$) to evaluate the registration accuracy. We use the number of parameters, GPU training memory in mebibytes (MiB), and the average GPU inference time in seconds (s) to measure model efficiency.

Implementation Details. Our method is implemented using PyTorch. Both the training and testing phases are deployed on an A100 GPU with 40GB VRAM. All models are optimized using Adam with a learning rate of 1×10^{-4} and batch size of 1 for 1000 epochs. On IXI, we employed the Normalized Cross-Correlation (NCC) loss with $\lambda = 2$. For 3D-CMR, we employed the Mean Squared Error (MSE) loss with $\lambda = 0.01$. C controls the model size, which is a trade-off between performance and efficiency. We used $C = 32$ for all experiments. We used Haar and Daubechies wavelet transform for the IXI and 3D-CMR datasets.

Comparative Methods. We compare our method with a series of state-of-the-art, including conventional SyN [2], U-Net-based methods (VoxelMorph (VM) [4], TransMorph (TM) [6]), model-driven methods (B-Spline-Diff [18], Fourier-Net [13]), and three pyramid methods (LapIRN [17], PR++ [15], ModeT [23]). All methods are trained with the official released code and the optimal parameters are tuned on the validation sets.

Quantitative and Qualitative Analysis. Table 1 shows the numerical results of compared methods on the IXI and 3D-CMR datasets. It can be seen that our method outperforms all compared methods in terms of Dice. For example, compared with the three pyramid methods, our method achieves at least 0.5% and 5.8% increase on IXI and 3D-CMR. Notably, Fig. 3 illustrates that WiNet requires significantly less GPU memory than pyramid methods in training, con-

suming only 31.9%, 25.6%, and 23.2% of the memory compared to LapIRN, PRNet++, and ModeT, respectively, and presenting faster speed than all three of them. Moreover, WiNet-Diff outperforms the iterative SyN with margins of 12% and 11% Dice on IXI and 3D-CMR, respectively. WiNet-Diff also surpasses model-driven approaches, achieving improvements of 0.4% and 0.5% (Fourier-Net-Diff) in Dice with 9% and 4% less memory usage for the IXI and 3D-CMR datasets, respectively, and securing the second-best performance in terms of HD.

As assessed in Fig. 4, WiNet-Diff can achieve 0% and 0.007% $|J|_{<0}$ % on IXI and 3D-CMR, respectively, resulting in a plausible deformation field compared to our WiNet. Additionally, we can observe that the estimated deformation fields from model-driven methods including B-Spline-Diff, Fourier-Net, and our WiNet tend to be smoother than U-Net-based methods (TM) and pyramid methods.

4 Conclusion

In this paper, we propose a model-driven WiNet to perform coarse-to-fine 3D medical image registration by estimating various frequency wavelet coefficients at different scales. Inheriting the properties of the wavelet transform, WiNet learns the low-frequency coefficients explicitly and high-frequency coefficients incrementally, making it explainable, accurate, and GPU efficient.

Acknowledgments. The research were performed using the Baskerville Tier 2 HPC service. Baskerville was funded by the EPSRC and UKRI through the World Class Labs scheme (EP/T022221/1) and the Digital Research Infrastructure programme (EP/W032244/1) and is operated by Advanced Research Computing at the University of Birmingham.

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

References

1. Ashburner, J.: A fast diffeomorphic image registration algorithm. *Neuroimage* **38**(1), 95–113 (2007)
2. Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C.: Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Med Image Anal* **12**(1), 26–41 (Feb 2008)
3. Avants, B.B., Tustison, N.J., Song, G., Cook, P.A., Klein, A., Gee, J.C.: A reproducible evaluation of ants similarity metric performance in brain image registration. *Neuroimage* **54**(3), 2033–2044 (2011)
4. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: Voxelmorph: a learning framework for deformable medical image registration. *IEEE transactions on medical imaging* **38**(8), 1788–1800 (2019)
5. Beg, M.F., Miller, M.L., Trounev, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International journal of computer vision* **61**, 139–157 (2005)

6. Chen, J., Frey, E.C., He, Y., Segars, W.P., Li, Y., Du, Y.: Transmorph: Transformer for unsupervised medical image registration. *Medical image analysis* **82**, 102615 (2022)
7. Chen, J., He, Y., Frey, E.C., Li, Y., Du, Y.: Vit-v-net: Vision transformer for unsupervised volumetric medical image registration. arXiv preprint arXiv:2104.06468 (2021)
8. Dalca, A., Balakrishnan, G., Guttag, J., Sabuncu, M.R.: Unsupervised learning of probabilistic diffeomorphic registration for images and surfaces. *Medical image analysis* **57**, 226–236 (2019)
9. Dalca, A.V., Balakrishnan, G., Guttag, J., Sabuncu, M.R.: Unsupervised learning for fast probabilistic diffeomorphic registration. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 729–738. Springer (2018)
10. De Vos, B.D., Berendsen, F.F., Viergever, M.A., Sokooti, H., Staring, M., Išgum, I.: A deep learning framework for unsupervised affine and deformable image registration. *Medical image analysis* **52**, 128–143 (2019)
11. Duan, J., Bello, G., Schlemper, J., Bai, W., Dawes, T.J., Biffi, C., de Marvao, A., Doumoud, G., O’Regan, D.P., Rueckert, D.: Automatic 3d bi-ventricular segmentation of cardiac images by a shape-refined multi-task deep learning approach. *IEEE transactions on medical imaging* **38**(9), 2151–2164 (2019)
12. Hu, B., Zhou, S., Xiong, Z., Wu, F.: Recursive decomposition network for deformable image registration. *IEEE Journal of Biomedical and Health Informatics* **26**(10), 5130–5141 (2022)
13. Jia, X., Bartlett, J., Chen, W., Song, S., Zhang, T., Cheng, X., Lu, W., Qiu, Z., Duan, J.: Fourier-net: Fast image registration with band-limited deformation. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 37, pp. 1015–1023 (2023)
14. Jia, X., Thorley, A., Chen, W., Qiu, H., Shen, L., Styles, I.B., Chang, H.J., Leonardis, A., De Marvao, A., O’Regan, D.P., et al.: Learning a model-driven variational network for deformable image registration. *IEEE Transactions on Medical Imaging* **41**(1), 199–212 (2021)
15. Kang, M., Hu, X., Huang, W., Scott, M.R., Reyes, M.: Dual-stream pyramid registration network. *Medical image analysis* **78**, 102379 (2022)
16. Kim, B., Kim, D.H., Park, S.H., Kim, J., Lee, J.G., Ye, J.C.: Cyclemorph: cycle consistent unsupervised deformable image registration. *Medical image analysis* **71**, 102036 (2021)
17. Mok, T.C., Chung, A.C.: Large deformation diffeomorphic image registration with laplacian pyramid networks. In: *Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part III* 23. pp. 211–221. Springer (2020)
18. Qiu, H., Qin, C., Schuh, A., Hammernik, K., Rueckert, D.: Learning diffeomorphic and modality-invariant registration using b-splines. In: *Medical Imaging with Deep Learning* (2021)
19. Sotiras, A., Davatzikos, C., Paragios, N.: Deformable medical image registration: A survey. *IEEE transactions on medical imaging* **32**(7), 1153–1190 (2013)
20. Stollnitz, E.J., DeRose, A.D., Salesin, D.H.: Wavelets for computer graphics: a primer. 1. *Ieee computer graphics and applications* **15**(3), 76–84 (1995)
21. Thorley, A., Jia, X., Chang, H.J., Liu, B., Bunting, K., Stoll, V., de Marvao, A., O’Regan, D.P., Gkoutos, G., Kotecha, D., et al.: Nesterov accelerated admm for fast diffeomorphic image registration. In: *Medical Image Computing and Computer*

- Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part IV 24. pp. 150–160. Springer (2021)
22. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Diffeomorphic demons: Efficient non-parametric image registration. *NeuroImage* **45**(1), S61–S72 (2009)
 23. Wang, H., Ni, Dong and Wang, Y.: Modet: Learning deformable image registration via motion decomposition transformer. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*. pp. 740–749 (2023)
 24. Yang, X., Wu, N., Cheng, G., Zhou, Z., David, S.Y., Beitler, J.J., Curran, W.J., Liu, T.: Automated segmentation of the parotid gland based on atlas registration and machine learning: a longitudinal mri study in head-and-neck radiation therapy. *International Journal of Radiation Oncology*Biography*Physics* **90**(5), 1225–1233 (2014)
 25. Zhang, J.: Inverse-consistent deep networks for unsupervised deformable image registration. *arXiv preprint arXiv:1809.03443* (2018)
 26. Zhao, S., Dong, Y., Chang, E.I., Xu, Y., et al.: Recursive cascaded networks for unsupervised medical image registration. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 10600–10610 (2019)
 27. Zhao, S., Lau, T., Luo, J., Eric, I., Chang, C., Xu, Y.: Unsupervised 3d end-to-end medical image registration with volume tweening network. *IEEE journal of biomedical and health informatics* **24**(5), 1394–1404 (2019)