



F3RNet: full-resolution residual registration network for deformable image registration

Zhe Xu^{1,2} · Jie Luo² · Jiangpeng Yan¹ · Xiu Li¹ · Jagadeesan Jayender²

Received: 2 September 2020 / Accepted: 24 March 2021 / Published online: 3 May 2021
© CARS 2021

Abstract

Purpose Deformable image registration (DIR) is essential for many image-guided therapies. Recently, deep learning approaches have gained substantial popularity and success in DIR. Most deep learning approaches use the so-called mono-stream high-to-low, low-to-high network structure and can achieve satisfactory overall registration results. However, accurate alignments for some severely deformed local regions, which are crucial for pinpointing surgical targets, are often overlooked. Consequently, these approaches are not sensitive to some hard-to-align regions, e.g., intra-patient registration of deformed liver lobes.

Methods We propose a novel unsupervised registration network, namely full-resolution residual registration network (F3RNet), for deformable registration of severely deformed organs. The proposed method combines two parallel processing streams in a residual learning fashion. One stream takes advantage of the full-resolution information that facilitates accurate voxel-level registration. The other stream learns the deep multi-scale residual representations to obtain robust recognition. We also factorize the 3D convolution to reduce the training parameters and enhance network efficiency.

Results We validate the proposed method on a clinically acquired intra-patient abdominal CT-MRI dataset and a public inspiratory and expiratory thorax CT dataset. Experiments on both multimodal and unimodal registration demonstrate promising results compared to state-of-the-art approaches.

Conclusion By combining the high-resolution information and multi-scale representations in a highly interactive residual learning fashion, the proposed F3RNet can achieve accurate overall and local registration. The run time for registering a pair of images is less than 3 s using a GPU. In future works, we will investigate how to cost-effectively process high-resolution information and fuse multi-scale representations.

Keywords Deformable image registration · Residual learning · Image-guided therapy · Deep learning

Introduction

In image-guided therapies (IGTs), e.g., preoperative planning, intervention and diagnosis, deformable image registration is the key to integrate complementary information contained in different time stamps or image modalities.

Therefore, developing fast and accurate deformable image registration methods is beneficial for the performance of IGT.

Traditional registration methods such as symmetric normalization (SyN) [1] align a pair of images by iteratively minimizing the appearance dissimilarity under regularization constraints. Furthermore, Deeds [12] utilizes discrete optimization, which shows promising results in abdominal registration [28]. However, solving a pairwise optimization is computationally intensive, resulting in slow speed in practice. Recently, due to the substantial improvement in computational efficiency over the traditional iterative registration, learning-based image registration approaches are becoming more prominent in task-specific and time-intensive applications [7]. Most learning-based registration approaches use fully supervised [3,4,20] or semi-supervised learning strategy [5,15] and heavily rely on ground-truth

✉ Xiu Li
li.xiu@sz.tsinghua.edu.cn
Jagadeesan Jayender
jayender@bwh.harvard.edu

¹ Shenzhen International Graduate School, Tsinghua University, Shenzhen 518055, China

² Brigham and Women's Hospital, Harvard Medical School, Boston 02115, USA

voxel correspondences and/or organ segmentation labels. Although these approaches struggle with imperfect ground-truth labels, they have made a significant impact on the field of deformable image registration. With the development of spatial transformer network (STN) [16], registration approaches that are based on unsupervised learning have also been introduced. For example, VoxelMorph [2] is a monomodal unsupervised registration framework that focuses on registering brain images of the same modality (unimodal registration). By modifying VoxelMorph, researchers have further proposed more unsupervised unimodal registration approaches [6,14,18,25].

Most existing learning-based registration approaches use the so-called mono-stream high-to-low, low-to-high network structure with augmented modules, e.g., skip-connection [2,8], multi-resolution fusion [14] and intermediate supervision [19]. This structure can significantly increase the size of the receptive field which is highly desirable for recognizing object information in images, but needs to recover the high-resolution information from the low-resolution representations. With increased receptive field sizes, these approaches prioritize overall registration accuracy, which is governed by the majority of easy-to-align regions, and overlook some severely deformed local regions. For example, livers with tumors usually have large local deformation due to progressed disease, and the deformations of the surrounding kidney and spleen are less significant. In a CT-to-MRI abdominal image registration, the aforementioned approaches are likely to estimate a deformation field that accurately registers kidney and spleen, yet perform poorly at local liver lobes alignment.

Besides, most of the image registration networks utilize 3D convolutional neural networks (3D CNNs) to exploit the semantic information in each CT/MRI slice and the spatial relationships across consecutive slices. It is understood that the training of 3D CNN is computationally expensive and may lead to insufficient training due to the small number of clinical datasets.

To address the above problems, we propose a novel unsupervised **full-resolution residual registration network (F3RNet)**, which is shown in Fig. 1(a). Distinct from the conventional mono-stream network structure, F3RNet consists of two parallel streams, namely “full-resolution stream” and “multi-scale residual stream.” Inspired by the success of using a high-resolution stream in human pose estimation and image inpainting tasks [9,23,26], “full-resolution stream” takes advantage of the detailed image information and facilitates accurate voxel-level registration, while the “multi-scale residual stream” learns the deep multi-scale residual representations to robustly recognize corresponding organs in both images and guarantee a high overall registration accuracy. Using the multi-scale residual block (MRB) modules, the network can progressively fuse information from the two

parallel streams in a residual learning fashion [10] to further boost the performance. In addition, we factorize the 3D convolution into two correlated 2D and 1D convolutions, thus effectively avoid over-parameterization [24].

To the best of our knowledge, we are the first to incorporate full-resolution representations with multi-scale high-level representations in a residual learning fashion to boost deformable image registration performance. The main contributions of our work can be summarized as follows:

- Our approach can unite the strong capability of capturing deep multi-scale representations with precise full-resolution spatial localization of the anatomical structures by interactively combining two parallel streams via the proposed MRB module and the residual learning mechanism. By taking into account such full-resolution information, the registration network is more sensitive to the hard-to-align regions and can provide better alignments for severely deformed local regions.
- The factorization of 3D convolution can markedly reduce the training parameters and enhance the network efficiency.
- We validate the proposed F3RNet on a clinically acquired intra-patient abdominal CT-MRI dataset and a public inspiratory and expiratory thoracic CT dataset. The experimental results on both multimodal and unimodal registration show that our method achieves superior performance over the existing state-of-the-art traditional and learning-based methods.

The outline of the paper is as follows: “Methods” section describes the details of our F3RNet, “Experiments” section presents the experimental details and registration results on both multimodal and unimodal datasets, and “Conclusions” section will draw conclusions of the paper.

Methods

Representing the moving image as I_m and the fixed image as I_f , medical image registration aims to estimate an optimal deformation field ϕ with three channels (x , y , z displacements) that can align I_m-I_f . In this section, we present our full-resolution residual registration network (shown in Fig. 1) firstly. Then, we describe the detailed structure of the designed residual block (RB) and multi-scale residual block (MRB), respectively. The factorization of 3D convolution is presented in “Factorized 3D convolution (F3D)” section, and the loss function of our network is described in “Loss function” section.

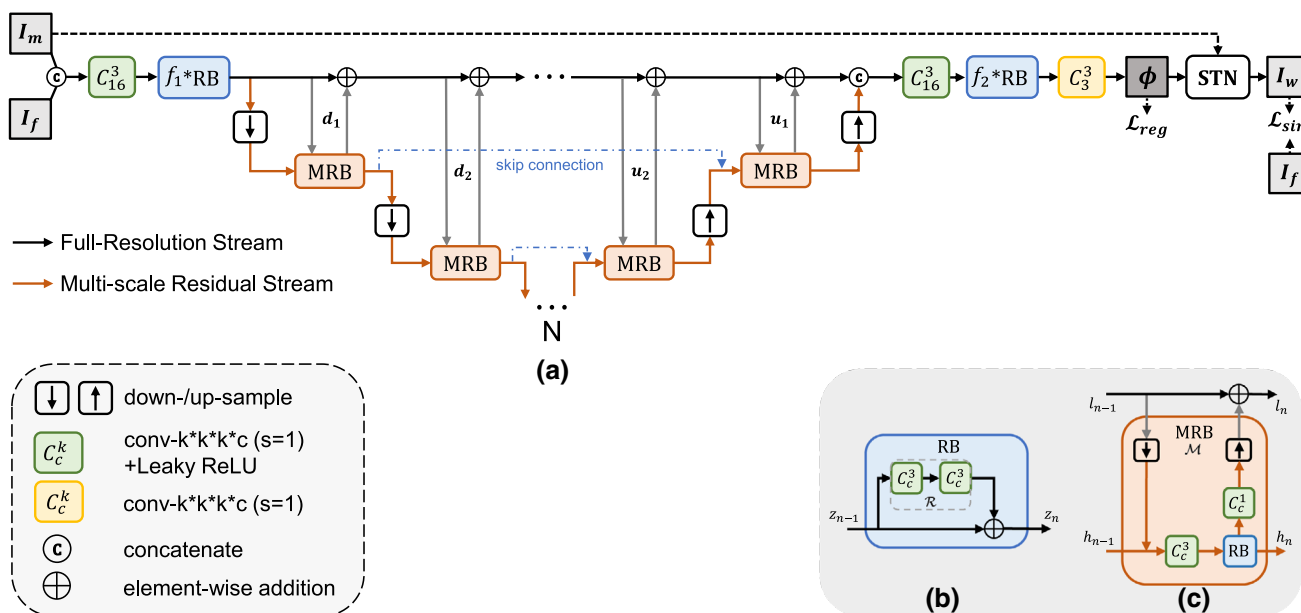


Fig. 1 Illustration of full-resolution residual stream network (F3RNet). **a** shows the overview of our F3RNet; **b** shows the residual block (RB); **c** shows the multi-resolution residual block (MRB). The network learns

parameters for a dense deformation field ϕ that aligns the moving image I_m to the fixed image I_f . N denotes the minimum volume is $(1/2^N)$ the size of the input images

Overview of the network

Distinct from the regular *high-to-low, low-to-high* one-pass network architecture, full-resolution residual registration network (F3RNet) unifies two parallel streams:

- **Full-resolution Stream.** Maintaining high-resolution features has demonstrated its superior performance for dense prediction [9,22,23,26]. The black line in Fig. 1a indicates the data flow of the full-resolution stream. This stream first concatenates I_m and I_f , followed by a 3D convolution and a series of residual blocks (RB, described in “Residual block (RB)” section). Then, the low-level features on this stream are successively computed by adding the residual from the other parallel stream. After that, the full-resolution stream reduces the number of channels via consecutive RBs and 3D convolutions step-by-step and estimates the 3-channel deformation field ϕ . Spatial transformation network (STN) [16] is applied to warp the moving image I_m with ϕ , so that the similarity between the warped image I_w and fixed image I_f can be evaluated. This stream does not employ any downsampling operation, resulting in good boundary localization but poor performance in deep semantic recognition. As such, the hard-to-align regions are propagated throughout the stream. Specifically, the convolutions in the full-resolution stream are all with 16 channels in our experiments except for the final 3-channel convolution used to generate the deformation field.

- **Multi-scale Residual Stream.** The data flow of multi-scale residual stream is depicted as the orange line in Fig. 1a. In contrast to the full-resolution stream, this stream is good at capturing high-level features that can improve the organ recognition performance. Specifically, successive pooling and convolution operations are leveraged to increase the receptive fields and enhance the robustness against small noises in the images. We also inherit the skip-connection design in regular *high-to-low, low-to-high* architecture that the feature spaces with same resolution are skip-connected by addition operation. Besides, with the help of our proposed multi-scale residual blocks (MRBs) that can simultaneously operate on both streams, the high-level features can directly interact with low-level features. The interior architecture of MRB is shown in Fig. 1c with elaboration in “Multi-scale residual block (MRB)” section. In our experiments, we set N to 4, which is the same as VoxelMorph [2], denoting that the lowest resolution is 1/16 of the original image. Specifically, at the resolution of 1/2 and 1/4 scale, the channel number of the feature map is set to 16. At the resolution of 1/8 and 1/16 scale, the number of feature channels becomes 32.

The information of the two distinct streams is automatically fused via residual learning [10]. By repeatedly fusing the features between two streams via computing successive multi-scale residuals, the full-resolution representations become rich for the dense deformation field prediction. At

the same time, richer low-level full-resolution information can in turn enhance the high-level multi-scale information.

Residual block (RB)

ResNets, proposed in [10], have demonstrated that residual learning can improve the training characteristics over traditional one-pass feed-forward learning. The interior architecture of the residual block (RB) is depicted in Fig. 1b. The output z_n of the RB can be formulated as:

$$z_n = z_{n-1} + \mathcal{R}(z_{n-1}), \quad (1)$$

where \mathcal{R} represents the residual branch consisting of two 3D convolutions with a kernel size of $3 \times 3 \times 3$ followed by LeakyReLU activations. Instead of computing z_n directly as in the traditional feed-forward network, the convolutional branch only needs to compute the residual \mathcal{R} in this architecture.

Multi-scale residual block (MRB)

The multi-scale residual block (MRB) follows the basic idea of residual block (RB) but elegantly achieves interaction between the full-resolution stream and multi-scale residual stream. An MRB consists of a series of pooling, 3D convolution and upsampling layers, as shown in Fig. 1c. Each MRB has two inputs, l_{n-1} as full-resolution low-level features and h_{n-1} as multi-resolution high-level features, and two corresponding outputs l_n and h_n . Intuitively, denoting the entire MRB operation as \mathcal{M} , the output l_n can be computed as:

$$l_n = l_{n-1} + \mathcal{M}(l_{n-1}, h_{n-1}). \quad (2)$$

Specifically, first, the resolution of l_{n-1} is reduced to that of h_{n-1} by a pooling operation, followed by a feature map concatenation. Then, the concatenated feature map undergoes a 3D convolution with a kernel size of $3 \times 3 \times 3$, followed by a Residual Block (RB) with the same number of channels, and the output h_n is connected to the next process of the multi-scale residual stream. Meanwhile, the output of the $3 \times 3 \times 3$ convolutional module adjusts the number of channels and the resolution to be consistent with l_{n-1} through a $1 \times 1 \times 1$ convolutional bottleneck layer and an upsampling layer at the other end. By such a process, we can readily use addition operations to integrate the residuals learned in the MRB in the full-resolution stream, thus forming a dual-stream highly interactive residual module.

Factorized 3D convolution (F3D)

Most medical images, as shown in Fig. 2a, consist of 3D image stacks with the size of $W \times H \times D$, where W , H ,




Fig. 2 Illustration of **a** 3D medical image scans, **b** regular 3D convolution with kernel size of $3 \times 3 \times 3$, **c** F3D convolution block

D represents the width, height, and the number of sequential slices. Inspired by the Inception [24] where large 2D convolution is factorized into two smaller ones, we factorize 3D convolution block for learning the volumetric representation. Specifically, suppose that we have a 3D convolution with kernel size of $3 \times 3 \times 3$ (Fig. 2b), it can be factorized into a $3 \times 3 \times 1$ convolution and a $1 \times 1 \times 3$ convolution in a cascaded fashion (Fig. 2c) to continuously capture dense 2D features in $W \times H$ slices with 1D attention weights that build sparse sequential relationships across adjacent slices. As such, the number of trainable parameters is reduced from $O(3^3 = 27)$ to $O(3 \times 3 + 3 = 12)$, where we can reduce the parameters by half.

However, it is noteworthy that the factorization is not totally equivalent to regular 3D convolution, and a further ablation study over factorized 3D convolution is presented in “Ablation study of F3D convolution” section.

Loss function

The loss function of our network consists of two components as shown in Eq. (3). The similarity loss \mathcal{L}_{sim} penalizes the dissimilarity between the fixed image I_f and the warped image $I_w(I_m \circ \phi)$. The deformation regularization \mathcal{L}_{reg} adopts an L2-norm of the gradients of the final deformation field ϕ with a trade-off weight λ . We write the total loss as:

$$\mathcal{L}(I_m, I_f, \phi) = \mathcal{L}_{\text{sim}}(I_f, I_m \circ \phi) + \lambda \mathcal{L}_{\text{reg}}(\phi). \quad (3)$$

Specifically, *modality independent neighborhood descriptor (MIND)* [11] can be used to measure the similarity of both multimodal and unimodal images. MIND is a modality-invariant structural representation, and we can minimize the difference in the MIND features between the warped image $I_w(I_m \circ \phi)$ and the fixed image I_f to effectively train the registration network. We define:

$$\mathcal{L}_{\text{sim}}(I_f, I_m \circ \phi) = \frac{1}{N|R|} \sum_x \| \text{MIND}(I_m \circ \phi) - \text{MIND}(I_f) \|_1, \quad (4)$$

where N denotes the number of voxels in input images $I_w(I_m \circ \phi)$ and I_f , R is a non-local region around voxel x .

Experiments

Dataset and implementation

In this work, we focus on the application of abdominal CT-MRI multimodal registration to improve the accuracy of percutaneous nephrolithotomy (PCNL). To further validate the effectiveness of our method, we also evaluate the proposed method on a public lung CT unimodal dataset [13].

- *Abdominal CT-MRI dataset*: Under the IRB approved study, we obtained an proprietary intra-patient CT-MRI dataset containing paired CT and MR images from 50 patients. The liver, kidney and spleen in both CT and MRI were manually segmented for quantitative evaluation. Standard preprocessing steps, including affine spatial normalization, resampling and intensity normalization, were performed. The images were cropped into $144 \times 144 \times 128$ subvolume with 1mm isotropic voxels and divided into two groups for training (40 cases) and testing (10 cases).
- *Learn2reg 2020 Lung CT dataset* [13]: This dataset contains paired inspiratory and expiratory thorax CT images from 30 subjects (20 cases for training and 10 cases for testing). For all scans, the lung segmentation masks are provided for evaluation. Standard preprocessing steps, including affine spatial normalization and resampling, had been performed by the challenge organization. We further carried out intensity normalization and cropped images into $128 \times 128 \times 160$ subvolume.

The proposed method is implemented using Keras with the Tensorflow backend. We train the network on a NVIDIA Titan X (Pascal) GPU using Adam optimizer [17] with a learning rate of $1e-5$. The batch size is set to 1. As for the optimal trade-off weight λ , we conduct exhaustive grid search and select the value that achieves the highest average Dice scores of ROIs on hold-out test set.

Measurement

We evaluate the registration performance using a series of metrics for each method, mainly including average surface distance (ASD) (lower is better) and the average Dice score (higher is better) between the segmentation masks of warped images and fixed images. Besides, the average number of voxels with non-positive Jacobian determinant ($|J_\phi| \leq 0$) in the deformation fields is counted for evaluating the dif-




Fig. 3 Results of varying the trade-off weight λ on average Dice score of ROIs

feomorphism of the local deformation (lower is better). The standard deviation of the Jacobian determinant ($\sigma(|J_\phi|)$) is also calculated for evaluating the smoothness of transformations (lower is better).

Experimental results

Ablation study of F3D convolution


As mentioned in “Factorized 3D convolution (F3D)” section, although convolution factorization can dramatically reduce the training parameters, it may not be totally equivalent to the regular 3D convolution in practice. Therefore, we investigate the different combinations of F3D convolution in our F3RNet. In our experiments, except for the final 3-channel 3D convolution used to generate the deformation field, other $3 \times 3 \times 3$ convolutions can be replaced. The variants of F3RNet are presented in Table 1. In particular, the number of parameters of F3RNet-w/ F3D is only 56.8% of the original F3RNet. “More MRBs” indicate that two extra MRBs are added at the lowest resolution path, which means that it is possible to use the reduced parameters to add more MRBs to enhance the network’s learning capability.

Figure 3 presents the average Dice scores of ROIs on the hold-out test set for varying values of the smoothing trade-off weight λ . The best Dice scores occur when $\lambda = 1.5$ for F3RNet-w/o F3D, F3RNet-w/ F3D, F3RNet-Dec, F3RNet-FR and F3RNet-MRB, and $\lambda = 2$ for F3RNet-Enc and F3RNet-MS. In particular, F3RNet-w/o F3D and F3RNet-MRB achieve better Dice scores than all other variants. Moreover, after achieving the best Dice scores at $\lambda = 1.5$, the results vary slowly over larger λ for F3RNet-w/o F3D and F3RNet-MRB, showing that the two models are more robust to the choice of λ .

Figure 4 shows visual results of warped images for the ablation analysis. We can firstly see that the original F3RNet (F3RNet-w/o F3D) can effectively register the multimodal

Table 1 Different combinations of F3D convolution (✓) in proposed F3RNet

Network	FR stream	MS stream		More MRBs
		Encoder	Decoder	
F3RNet-w/o F3D				
F3RNet-w/ F3D	✓	✓	✓	
F3RNet-Enc		✓		
F3RNet-Dec			✓	
F3RNet-FR	✓			
F3RNet-MS		✓	✓	
F3RNet-MRB	✓	✓	✓	✓

**Fig. 4** Visual results of an example for CT-to-MRI registration. Outside the grey box shows an example fixed MR image and a zoom-in region with the segmentation masks of the liver (green), kidney (red), and spleen (blue). The corresponding warped CT images and zoom-in

regions for baselines and ablation study are presented in the grey box. A good registration will cause structures in warped images to close to the corresponding fixed segmentation masks. The red arrows indicate the registration of interest at the boundary of the organ

images. If we replace all 3D convolutions with F3D (F3RNet-w/ F3D) or only replace the convolution in encoder and decoder (F3RNet-Enc and F3RNet-Dec), our methods can still effectively register the CT image but have slight performance degradation. Interestingly, if we replace the regular convolution on the entire multi-scale residual stream or full-resolution stream alone, this will cause the information of the two streams to not effectively interact and introduce noise, resulting in unstable performance and significant registration degradation. Therefore, if we use F3D to reduce the model parameters, the 3D convolution on both streams should be

replaced at the same time. Further, we can use the reduced parameters to add more MRBs (F3RNet-MRB). From the visual results, it can be seen that the registration performance is either maintained or slightly improved.

Table 2 also provides the comprehensive quantitative results for all baseline methods and the variants of our F3RNet with different combinations of F3D. As for the results for ablation analysis, we can see that F3RNet-w/o F3D and F3RNet-MRB achieve the best performance. Specifically, with only 80.2% parameters of F3RNet-w/o F3D, F3RNet-MRB achieves better ASD results in the liver and

Table 2 Average Dice scores and average ASD evaluations (mean ± std) for CT-to-MRI registration of all baseline methods and F3RNet with different combinations of F3D

Methods	ASD(mm)		Dice(%)		Kidney		Spleen		Kidney	Spleen	Times(s)	GPU/CPU
	Liver	Kidney	Liver	Kidney	Liver	Kidney	Liver	Kidney				
Moving	4.95±0.82	1.97±0.52	2.01±0.36	77.18±4.13	78.24±3.21	80.14±3.17	-	-	-	-	-	-
SyN	4.81±0.79	1.54±0.63	1.92±0.41	79.18±4.37	80.21±3.41	82.91±3.08	137.60(0.41)	-	-	-	-97	-
Deeds	3.02±0.51	1.30±0.61	1.35±0.42	85.27±3.37	83.86±2.33	83.37±3.19	4.80(0.27)	-	-	-	-37	-
VM-1	3.98±0.74	1.42±0.54	1.75±0.52	82.39±4.11	82.83±2.68	82.34±2.97	0.30(0.11)	0.260	0.300	0.300	1.32/21	1.33/23
VM-2	3.92±0.53	1.47±0.37	1.72±0.49	84.17±3.57	82.76±2.98	83.51±3.36	1.10(0.15)	0.217	0.217	0.217	1.29/20	1.29/20
FAIM	3.88±0.73	1.51±0.63	1.66±0.41	84.51±3.76	81.99±2.84	83.12±3.24	0.00(0.12)	0.359	0.359	0.359	2.31/26	2.31/26
F3RNet-w/o F3D	2.19±0.37	1.28±0.39	1.37±0.38	86.65±3.42	85.39±2.75	83.58±3.18	0.40(0.11)	0.204	0.204	0.204	1.22/25	1.22/25
F3RNet-w/ F3D	2.63±0.52	1.36±0.86	1.41±0.72	85.16±4.19	84.39±2.23	82.98±3.43	0.80(0.09)	0.335	0.335	0.335	1.93/25	1.93/25
F3RNet-Dec	2.77±0.69	1.33±0.69	1.43±0.48	85.87±4.23	84.18±2.77	83.06±3.92	0.20(0.19)	0.321	0.321	0.321	1.79/23	1.79/23
F3RNet-Enc	2.73±0.47	1.39±0.52	1.38±0.36	85.32±4.35	84.27±3.35	83.11±3.28	0.50(0.17)	0.348	0.348	0.348	2.13/25	2.13/25
F3RNet-FR	3.82±0.59	1.53±0.39	1.56±0.43	81.43±4.27	82.97±3.18	83.13±3.52	0.70(0.24)	0.243	0.243	0.243	1.41/24	1.41/24
F3RNet-MS	3.94±0.75	1.51±0.44	1.52±0.67	82.31±3.84	83.06±3.33	82.99±3.28	1.20(0.24)	0.288	0.288	0.288	1.42/25	1.42/25
F3RNet-MRB	2.17±0.46	1.29±0.48	1.34±0.26	86.79±3.18	85.42±2.98	83.16±3.43	0.40(0.13)					

Best results are shown in bold. Besides, $|J_{\phi}| \leq 0$ (average number of folding voxels in the deformation fields) and $\sigma(|J_{\phi}|)$ (the smoothness of the deformation fields) are also provided

kidney registration than F3RNet-w/o F3D, while it also achieves better Dice score in liver and spleen registration with reasonable diffeomorphism and smoothness of the deformation fields. Meanwhile, consistent with the visual assessment, we can also see that F3RNet-FR and F3RNet-MS both yield significant performance degradation over ASD and Dice score as they cause the features of the two streams to be disjointed.

Comparison with baselines on abdominal CT-to-MRI registration

To evaluate our proposed method, five open-source state-of-the-art baseline approaches are also compared, including two traditional methods SyN [1] with mutual information (MI) metric [27] and Deeds [12] with five-levels of discrete optimization, and three unsupervised learning-based methods, marked as VoxelMorph-1 (VM-1) [2], VoxelMorph-2 (VM-2) [2], and FAIM [18]. The three learning-based methods are initially proposed for unimodal registration, and we extend them for both multimodal and unimodal registration by using MIND-based similarity metric. We use the same test set to search for the best regularization weights and then set the weights to 1.5 for VM-1, VM-2 and FAIM. Other parameters, such as learning rate and batch size, remain the same as our method.

Figure 4 also illustrates the warped CT images produced by other baseline methods. As we have mentioned above, liver registration is much more challenging in the abdominal image registration task. From the results, we can see that the traditional method SyN fails to align the liver with large local deformation while Deeds performs much better. As for other deep learning methods, VM-1, VM-2, and FAIM achieve relatively satisfactory performance but still have considerable disagreements. Except for F3RNet-FR and F3RNet-MS, our methods have the most visually appealing boundary alignment, which demonstrates that our F3RNet can better register the hard-to-align regions.

The quantitative results for the baseline methods are also presented in Table 2. Consistent with the visual results, the evaluations over ASD and Dice scores of our proposed methods are better than the traditional methods and other state-of-the-art unsupervised registration methods with reasonable quality of the deformation fields, except for F3RNet-FR and F3RNet-MS. Among the baseline methods, Deeds provides competitive results over SyN and other learning-based methods. Furthermore, the traditional methods take much more time (97s for SyN and 37s for Deeds) to register an image pair. In contrast, all deep learning methods can complete a registration task in 3 seconds with a GPU, making it appealing for image-guided therapies with intense time demand.




Fig. 5 Visual results of an example for MRI-to-CT registration. Outside the grey box shows an example fixed CT image and a zoom-in region with the segmentation masks of the liver (green), kidney (red),

and spleen (blue). The corresponding warped MR images and zoom-in regions for all methods are presented in the grey box. The red arrows indicate the registration of interest at the boundary of the organ

Table 3 Average Dice scores and average ASD evaluations (mean ± std) for MRI-to-CT registration

Methods	ASD(mm)			Dice(%)			$ J_\phi \leq 0$ ($\sigma(J_\phi)$)
	Liver	Spleen	Kidney	Liver	Spleen	Kidney	
Moving	4.95±0.82	1.97±0.52	2.01±0.36	77.18±4.13	78.24±3.21	80.14±3.17	–
SyN	4.73±0.68	1.63±0.57	1.99±0.33	78.36±4.67	78.53±3.39	81.38±2.98	89.2(0.33)
Deeds	3.58±0.53	1.56±0.44	1.84±0.40	83.59±3.87	82.63±3.41	83.06±3.49	1.20(0.14)
VM-1	4.02±0.73	1.61±0.64	1.95±0.28	81.34±4.06	80.72±3.02	82.46±3.07	0.10(0.05)
VM-2	3.59±0.67	1.53±0.52	1.87±0.36	83.28±4.03	82.81±3.14	83.37±2.83	0.00(0.03)
FAIM	3.71±0.87	1.51±0.63	1.89±0.31	84.33±3.64	81.06±3.48	83.44±2.92	0.00(0.04)
F3RNet-w/o F3D	3.12±0.59	1.43±0.59	1.68±0.42	85.75±4.11	83.02±3.26	84.07±3.04	0.00(0.06)
F3RNet-MRB	3.04±0.65	1.38±0.51	1.67±0.35	85.93±3.52	83.47±3.51	84.39±2.77	0.00(0.04)

Best results are shown in bold. Besides, $|J_\phi| \leq 0$ (average number of folding voxels in the deformation fields) and $\sigma(|J_\phi|)$ (the smoothness of the deformation fields) are also provided

Experiments on abdominal MRI-to-CT registration

Among all the proposed networks for CT-to-MRI registration, F3RNet-w/o F3D and F3RNet-MRB provide superior results. To further validate the effectiveness of the two proposed methods, we also perform the MRI-to-CT registration in turn. The division of the dataset and the other training settings of the networks, e.g., regularization trade-off weights, etc., are consistent with the CT-to-MRI registration task.

The visualization of the registration results in Fig. 5 shows that our methods, F3RNet-w/o F3D and F3RNet-MRB, achieve more accurate organ alignment than other traditional and deep learning approaches, especially for the liver.

The quantitative evaluation of MRI-to-CT registration is summarized in Table 3. Our proposed methods achieve better results in terms of ASD and Dice scores than that of the traditional method and other state-of-the-art unsupervised learning registration methods. In particular, F3RNet-MRB achieves the best registration accuracy among all the methods with reasonably low $|J_\phi| \leq 0$ and $\sigma(|J_\phi|)$.

Experiments on expiration-to-inspiration lung CT registration

Apart from the large local deformation between expiratory and inspiratory lung CT images, another challenge of the Learn2Reg 2020 Lung CT dataset [13] is that the lungs are not fully visible in several expiratory scans as shown in I_m in Fig. 6. In our experiment, MIND-based similarity metric [11] is still used to guide the network training. Empirically, the regularization weights are all set to 1.5 for VM-1, VM-2, FAIM and F3RNet. Other parameters, such as learning rate and batch size, remain the same as the aforementioned experiments.

We visualize an example of the registration results from both axial and coronal views in Fig. 6. Apparently, the proposed methods, F3RNet-w/o F3D and F3RNet-MRB, achieve more accurate lung alignment than other traditional and deep learning approaches, especially from the coronal view.

The quantitative evaluation of expiration-to-inspiration lung CT registration is summarized in Table 4. Our proposed methods achieve better results in terms of ASD and Dice




Fig. 6 Visual results of an example for expiration-to-inspiration lung CT registration from both axial and coronal views. The red contours represent the lung segmentation of the fixed inspiratory CT image

Table 4 Average Dice scores and average ASD evaluations (mean \pm std) for lung CT registration

Methods	ASD(mm)	Dice(%)	$ J_\phi \leq 0$ ($\sigma(J_\phi)$)
Moving	2.51 \pm 0.83	86.37 \pm 3.45	–
SyN(MI)	2.34 \pm 0.31	86.27 \pm 2.62	133.50(0.36)
Deeds	2.35 \pm 0.72	87.28 \pm 2.78	7.50(0.11)
VM-1	2.46 \pm 0.84	86.63 \pm 3.49	0.00(0.08)
VM-2	2.38 \pm 0.82	87.05 \pm 3.39	1.20(0.09)
FAIM	2.43 \pm 0.85	86.86 \pm 3.39	4.50(0.11)
F3RNet-w/o F3D	1.95 \pm 0.37	87.84 \pm 2.50	1.80(0.16)
F3RNet-MRB	1.69\pm0.34	88.95\pm2.59	2.10(0.13)

Best results are shown in bold. Besides, $|J_\phi| \leq 0$ (average number of folding voxels in the deformation fields) and $\sigma(|J_\phi|)$ (the smoothness of the deformation fields) are also provided

scores than that of the traditional methods and other state-of-the-art unsupervised learning registration networks with reasonable tradeoff in the diffeomorphism and smoothness of the deformation fields. In particular, F3RNet-MRB achieves the best performance among all the methods.

Conclusions

In this work, we propose a novel unsupervised registration network, namely full-resolution residual registration network (F3RNet), which takes advantage of full-resolution information, multi-scale fusion, deep residual learning framework and 3D convolution factorization, to improve the deformable registration performance. The experimental results on both multimodal and unimodal tasks indicate that our network can better register the hard-to-align region, yielding superior accuracy of registration. In our experiments, we found the current input size to be a compromise between image resolution and GPU memory limitation. Recently, the Laplacian pyramid image registration network (LapIRN) [21] that includes three pyramid branches to register the image pairs at different resolutions with a coarse-to-fine optimization scheme is proposed, which brings promising enlightenment on improving multi-scale fusion-based registration. Future works will continuously focus on the lighter and more elegant ways to leverage high-resolution information and multi-scale

fusion to cope with the large local deformation under limited GPU memory.

Acknowledgements This project was supported by the National Institutes of Health (Grant Nos. R01EB025964, R01DK119269, and P41EB015898), the National Key R&D Program of China (No. 2020AAA0108303), NSFC 41876098 and the Overseas Cooperation Research Fund of Tsinghua Shenzhen International Graduate School (Grant No. HW2018008).

Declarations

Conflict of interest Jayender Jagadeesan owns equity in Navigation Sciences, Inc. He is a co-inventor of a navigation device to assist surgeons in tumor excision that is licensed to Navigation Sciences. Dr. Jagadeesan's interests were reviewed and are managed by BWH and Partners HealthCare in accordance with their conflict of interest policies.

Ethical approval All procedures performed in studies involving human participants were in accordance with the ethical standards of the institutional and/or national research committee and with the 1964 Helsinki declaration and its later amendments or comparable ethical standards.

Informed consent Informed consent was obtained from all individual participants included in the study.

References

1. Avants BB, Epstein CL, Grossman M, Gee JC (2008) Symmetric diffeomorphic image registration with cross-correlation: evaluating

- automated labeling of elderly and neurodegenerative brain. *Med Image Anal* 12(1):26–41
2. Balakrishnan G, Zhao A, Sabuncu MR, Guttag J, Dalca AV (2018) An unsupervised learning model for deformable medical image registration. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 9252–9260
 3. Chee E, Wu Z (2018) Airnet: self-supervised affine registration for 3d medical images using neural networks. *arXiv preprint arXiv:1810.02583*
 4. Fan J, Cao X, Yap PT, Shen D (2019) Birnet: brain image registration using dual-supervised fully convolutional networks. *Med Image Anal* 54:193–206
 5. Ferrante E, Dokania PK, Silva RM, Paragios N (2018) Weakly supervised learning of metric aggregations for deformable image registration. *IEEE J Biomed Health Inform* 23(4):1374–1384
 6. Ferrante E, Oktay O, Glocker B, Milone DH (2018) On the adaptability of unsupervised cnn-based deformable image registration to unseen image domains. In: *International workshop on machine learning in medical imaging*. Springer, pp. 294–302
 7. Fu Y, Lei Y, Wang T, Curran WJ, Liu T, Yang X (2020) Deep learning in medical image registration: a review. *Phys Med Biol* 65(20):20TR01
 8. Ghosal S, Ray N (2017) Deep deformable registration: enhancing accuracy by fully convolutional neural net. *Pattern Recognit Lett* 94:81–86
 9. Guo Z, Chen Z, Yu T, Chen J, Liu S (2019) Progressive image inpainting with full-resolution residual network. In: *Proceedings of the 27th acm international conference on multimedia*, pp. 2496–2504
 10. He K, Zhang X, Ren S, Sun J (2016) Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778
 11. Heinrich MP, Jenkinson M, Bhushan M, Matin T, Gleeson FV, Brady M, Schnabel JA (2012) Mind: modality independent neighbourhood descriptor for multi-modal deformable registration. *Med Image Anal* 16(7):1423–1435
 12. Heinrich MP, Jenkinson M, Brady M, Schnabel JA (2013) MRF-based deformable registration and ventilation estimation of lung CT. *IEEE Trans Med Imaging* 32(7):1239–1248
 13. Hering A, Murphy K, van Ginneken B (2020) Lean2Reg challenge: CT lung registration-training data. <https://doi.org/10.5281/zenodo.3835682>
 14. Hu X, Kang M, Huang W, Scott MR, Wiest R, Reyes M (2019) Dual-stream pyramid registration network. In: *International conference on medical image computing and computer-assisted intervention*. Springer, pp. 382–390
 15. Hu Y, Modat M, Gibson E, Li W, Ghavami N, Bonmati E, Wang G, Bandula S, Moore CM, Emberton M, Ourselin S, Noble JA, Barratt DC, Vercauteren T (2018) Weakly-supervised convolutional neural networks for multimodal image registration. *Med Image Anal* 49:1–13
 16. Jaderberg M, Simonyan K, Zisserman A, Kavukcuoglu K (2015) Spatial transformer networks, pp 2017–2025. *arXiv preprint arXiv:1506.02025*
 17. Kingma DP, Ba J (2015) Adam: A method for stochastic optimization. In: *International conference on learning representations*
 18. Kuang D, Schmäh T (2019) Faim—a convnet method for unsupervised 3d medical image registration. In: *international workshop on machine learning in medical imaging*. Springer, pp. 646–654
 19. Li H, Fan Y (2018) Non-rigid image registration using self-supervised fully convolutional networks without training data. In: *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, pp. 1075–1078. IEEE
 20. Lv J, Yang M, Zhang J, Wang X (2018) Respiratory motion correction for free-breathing 3d abdominal MRI using CNN-based image registration: a feasibility study. *Br J Radiol* 91:20170788
 21. Mok TC, Chung AC (2020) Large deformation diffeomorphic image registration with Laplacian pyramid networks. In: *International conference on medical image computing and computer-assisted intervention*. Springer, pp. 211–221
 22. Pohlen T, Hermans A, Mathias M, Leibe B (2017) Full-resolution residual networks for semantic segmentation in street scenes. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4151–4160
 23. Sun K, Xiao B, Liu D, Wang J (2019) Deep high-resolution representation learning for human pose estimation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5693–5703
 24. Szegedy C, Vanhoucke V, Ioffe S, Shlens J, Wojna Z (2016) Rethinking the inception architecture for computer vision. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826
 25. de Vos BD, Berendsen FF, Viergever MA, Sokooti H, Staring M, Išgum I (2019) A deep learning framework for unsupervised affine and deformable image registration. *Med Image Anal* 52:128–143
 26. Wang J, Sun K, Cheng T, Jiang B, Deng C, Zhao Y, Liu D, Mu Y, Tan M, Wang X, Liu W, Xiao B (2020) Deep high-resolution representation learning for visual recognition. *IEEE Trans Pattern Anal Mach Intell*. <https://doi.org/10.1109/TPAMI.2020.2983686>
 27. Wells WM III, Viola P, Atsumi H, Nakajima S, Kikinis R (1996) Multi-modal volume registration by maximization of mutual information. *Med Image Anal* 1(1):35–51
 28. Xu Z, Lee CP, Heinrich MP, Modat M, Rueckert D, Ourselin S, Abramson RG, Landman BA (2016) Evaluation of six registration methods for the human abdomen on clinically acquired CT. *IEEE Trans Biomed Eng* 63(8):1563–1572

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