

Dual-Stream Pyramid Registration Network

Xiaojun Hu^{*1,2}, Miao Kang^{*1,2}, Weilin Huang^{**1,2}, Matthew R. Scott^{1,2},
Roland Wiest³, and Mauricio Reyes⁴

¹ Malong Technologies, Shenzhen, China

² Shenzhen Malong Artificial Intelligence Research Center, Shenzhen, China

³ Support Center for Advanced Neuroimaging, Univ. Institute for Diagnostic and
Interventional Neuroradiology, Univ. Hospital, Bern, Switzerland

⁴ Insel Data Science Center, Inselspital, Univ. Hospital, Univ. of Bern, Switzerland

Abstract. We propose a Dual-Stream Pyramid Registration Network (referred as Dual-PRNet) for unsupervised 3D medical image registration. Unlike recent CNN-based registration approaches, such as VoxelMorph, which explores a single-stream encoder-decoder network to compute a registration field from a pair of 3D volumes, we design a two-stream architecture able to compute multi-scale registration fields from convolutional feature pyramids. Our contributions are two-fold: (i) we design a two-stream 3D encoder-decoder network which computes two convolutional feature pyramids separately for a pair of input volumes, resulting in strong deep representations that are meaningful for deformation estimation; (ii) we propose a pyramid registration module able to predict multi-scale registration fields directly from the decoding feature pyramids. This allows it to refine the registration fields gradually in a coarse-to-fine manner *via* sequential warping, and enable the model with the capability for handling significant deformations between two volumes, such as large displacements in spatial domain or slice space. The proposed Dual-PRNet is evaluated on two standard benchmarks for brain MRI registration, where it outperforms the state-of-the-art approaches by a large margin, e.g., having improvements over recent VoxelMorph [2] with $0.683 \rightarrow 0.778$ on the LPBA40, and $0.511 \rightarrow 0.631$ on the Mindboggle101, in term of average Dice score.

Keywords: Medical image registration, encoder-decoder network, deformable registration, and brain MRI.

1 Introduction

Deformable image registration has been widely used in image diagnostics, disease monitoring, and surgical navigation, with the goal of learning the anatomical correspondence between a moving image and a fixed image. A registration process mainly consists of three steps: establishing a deformation model, designing

* Authors contributed equally

** Weilin Huang is the corresponding author (e-mail:whuang@malong.com).

a similarity measurement function, and a parameter optimization step. Traditional deformable registration methods often cast it into a complex optimization problem that involves intensive computation by densely computing voxel-level similarities. Recent deep learning technologies have advanced this task significantly by developing learning-based approaches, which allow them to leverage strong feature learning capability of deep neural networks [2,3], resulting in fast training and inference, e.g., by just taking orders of magnitude less time [2].

Medical image registration often requires strong supervision information, such as ground-truth registration fields or anatomical landmarks. However, obtaining a large-scale medical dataset with such strong annotations is extremely expensive, which inevitably limits the applications of supervised approaches. Recently, unsupervised learning-based registration methods have been developed, by learning a registration function that maximizes the similarity between a moving image and a fixed image. For example, Balakrishnan *et al.* [2] proposed VoxelMorph which learns a parameterized registration function using a convolutional neural network (CNN). VoxelMorph estimates a deformation field by using an encoder-decoder CNN, and warps a moving image with a spatial transformation layer [1]. Kuang and Schmah [3] developed an unsupervised method, named as FAIM, which extends VoxelMorph by improving the network design with a new registration function. However, Lewis *et al.* [5] demonstrated that the performance of existing CNN-based approaches can be limited on challenging clinical applications where two medical images or volumes may have significant spatial displacements or large slice spaces. It is straightforward to handle such deformations by aligning them in a coarse-to-fine manner.

Contributions. Recent approaches on optical flow estimation [6] attempted to handle large displacements by gradually refining the estimated flows, which inspired the current work. They designed a cascaded architecture where multiple FlowNets were employed to gradually warp the images, while current paper describes a dual-stream single-model approach that implements sequential layer-wise refinements of registration fields, which in turn are used to warp the convolutional features rather than the images. This results in an end-to-end trainable model for 3D medical image registration.

Our work extends recent VoxelMorph [2] with two technical improvements, which are the key to boost the performance. 1) we develop a dual-stream encoder-decoder network able to compute two convolutional feature pyramids separately from a pair of input volumes, resulting in stronger deep representations for estimating multi-level deformations. 2) we propose a pyramid registration module which directly estimates registration fields from convolutional features, resulting in a set of layer-wise registration fields that encode multi-level context information by sequentially warping the convolutional features. This enables the model with the ability to work on significant deformations. 3) we evaluate our model on the LPBA40 and the Mindboggle101, where our method improved the results of VoxelMorph [2] considerably, with $0.683 \rightarrow 0.778$ and $0.511 \rightarrow 0.631$ respectively, in term of average Dice score.

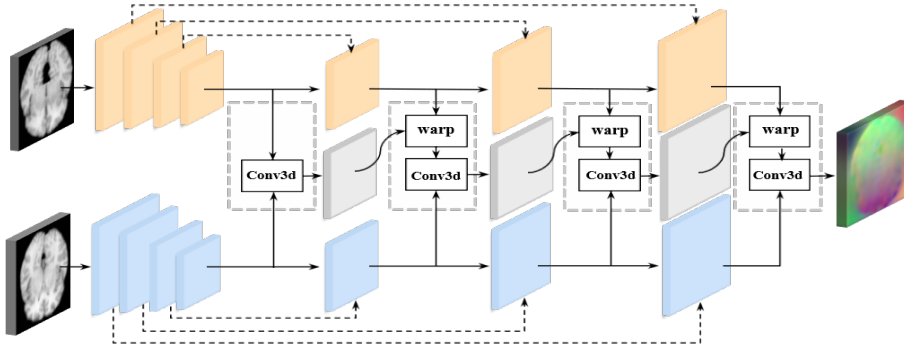


Fig. 1. Architecture of the proposed Dual-PRNet, which is a dual-stream encoder-decoder network, integrated with a new pyramid registration module.

2 Dual-Stream Pyramid Registration Network

In this section, we describe the details of the proposed Dual-PRNet, including two main components: (i) a dual-stream encoder-decoder network for computing feature pyramids, and (ii) a pyramid registration module that estimates layer-wise registration fields in the decoding process.

2.1 Preliminaries

The goal of 3D medical image registration is to estimate a deformation field Φ which can warp a moving volume $M \subset R^3$ to a fixed volume $F \subset R^3$, so that the warped volume $W = M \circ \Phi \subset R^3$ can be accurately aligned to the fixed one F . We use $M \circ \Phi$ to denote the application of a deformation field Φ to the moving volume with a warping operation, and image registration can be formulated as an optimization problem:

$$\hat{\Phi} = \arg \min_{\Phi} \mathcal{L}(F, M, \Phi), \quad \mathcal{L}(F, M, \Phi) = \mathcal{L}_{sim}(F, M \circ \Phi) + \lambda \mathcal{L}_{smooth}(\Phi) \quad (1)$$

where \mathcal{L}_{sim} is a function measuring image similarity between $M \circ \Phi$ and F , and \mathcal{L}_{smooth} is a regularization constraint on Φ , which enforces spatial smoothness. Both \mathcal{L}_{sim} and \mathcal{L}_{smooth} can be defined in various forms. For example, VoxelMorph [2] uses a CNN to compute a deformation field, $\Phi = f_{\theta}(F, M)$, where θ are learnable parameters of the CNN. The deformation warping is implemented by using a spatial transformer network [1], $M \circ \Phi = f_{stn}(M, \Phi)$. VoxelMorph uses a single-stream encoder-decoder architecture with skip connections, which are similar to U-Net [7]. A pair of volumes, M and F , are stacked as the input of VoxelMorph. More details of VoxelMorph are described in [2].

2.2 Dual-Stream Architecture

Our Dual-PRNet is built on the encoder-decoder architecture of VoxelMorph, but improves it by introducing a dual-stream design, as shown in Fig. 1. Specifically,

the backbone of Dual-PRNet consists of a dual-stream encoder-decoder with shared parameters. We adopt the encoder as the same architecture of VoxResNet [8], which contains four down-sampling convolutional blocks. Each block has a 3D down-sampling convolutional layer with stride of 2. Thus the encoder reduces the spatial resolution of input volumes by a factor of 16 in total. Except for the first block, the down-sampling convolutional layer is followed by two ResBlocks, each of which contains two convolutional layers with a residual connection, similar to ResNet [9]. Batch normalization (BN) and ReLU operations are applied. More details of the network are presented in Fig. 1 of Supplementary Material (SM).

In the decoder stage, we apply skip connections on the corresponding convolutional maps in the encoding and decoding process. The features are fused using a Refine Unit, where the lower-resolution convolutional maps are up-sampled and added into the higher-resolution ones, by using a $1 \times 1 \times 1$ convolution layer. Finally, we obtain two feature pyramids with multi-resolution convolutional features computed from the moving volume and the fixed volume separately.

Dual-stream design allows us to first compute meaningful feature pyramids from two input volumes, and then predict deformable fields from the learned, stronger and more discriminative convolutional features, which is the key to boost the performance. This is different from existing single-stream networks, such as [2] and [3], which compute the convolutional features from two stacked volumes, and jointly estimate deformation fields using same convolutional filters. Furthermore, our dual-stream architecture can generate two paired feature pyramids where layer-wise deformation fields can be computed at multiple scales, allowing the model to generate more meaningful deformation fields by designing a new pyramid registration module, which is described next.

2.3 Pyramid Registration Module

VoxelMorph computes a single deformation field from the convolutional features at the last up-sampling layer in the decoding process, which limits its capability for handling large-scale deformations. Our pyramid registration module is able to predict multiple deformation fields with different resolutions, and generate pyramid deformation fields. As shown in Fig. 1, it computes a deformation field from a pair of convolutional features at each decoding layer. Each deformation field is computed by using a sequence of operations with feature warping, stacking, and convolution, except for the first deformation field where feature warping is not implemented. This results in a sequence of deformation fields with increasing resolutions, starting from the lowest-resolution decoding layer to the highest-resolution one. Our network includes four decoding layers, and thus generates four deformation fields sequentially.

Specifically, the first deformation field (Φ_1) is computed at the first decoding layer. We stack the convolutional features from two feature pyramids, and apply a 3D convolution with size of $3 \times 3 \times 3$ to estimate the deformation field, which is a 3D volume with the same shape of the corresponding convolutional maps. This deformation field is able to extract coarse-level context information, such as

high-level anatomical structure of brain, which is then encoded into generating subsequent deformation field by feature warping: (i) the current deformation field is up-sampled by using bilinear interpolation with a factor of 2, denoted as $u(\Phi_1)$, and (ii) then is applied to warp the convolutional maps of the next layer from the moving volume, by using a grid sample operation, as shown in Fig. 1. Then the warped convolutional maps are stacked again with the corresponding convolutional features generated from the fixed volume, followed by a convolution operation to generate a new deformation field. This process is repeated at each decoding layer, and can be formulated as,

$$\Phi_i = C_i^{3 \times 3 \times 3} (P_i^M \circ u(\Phi_{i-1}), P_i^F) \quad (2)$$

where $i = 1, 2, \dots, N$, and is empirically set to 4, indicating four decoding layers. $C_i^{3 \times 3 \times 3}$ denotes a 3D convolution at the i -th decoding layer. \circ is the warping operation that maps the coordinates of P_i^M to P_i^F using $u(\Phi_{i-1})$, where P_i^M and P_i^F are the convolutional feature pyramids computed from the moving volume and the fixed volume at the i -th decoding layer.

Finally, the estimated deformation fields are warped sequentially and recurrently with up-sampling, to generate a final deformation field, which encodes meaningful multi-level context information with multi-scale deformations. This allows the model to propagate strong context information over hierarchical decoding layers, where the estimated deformation fields are refined gradually in a coarse-to-fine manner, and thus aggregate both high-level context information and low-level detailed features. The high-level context information enables our model with the ability to work on large-scale deformations, while the fine-scale features allows it to preserve detailed anatomical structure. We integrate pyramid registration module into the dual-stream architecture, resulting in an end-to-end trainable model. We adopt a negative local cross correlation (NLCC) as loss function, which is coupled with a smooth regularization, by simply following VoxelMorph [2].

3 Experimental Results and Comparisons

The proposed Dual-PRNet is evaluated on 3D brain MRI registration on two public datasets, LPBA40 [10] and Mindboggle101 [11].

Datasets. LPBA40 [10] contains 40 T1-weighted MR images, each of which was annotated with 56 subcortical ROIs. Mindboggle101 [11] contains 101 T1-weighted MR images, which were annotated with 25 cortical regions, and can be used to evaluate registration results regarding fine brain structure.

Experimental Settings. Experiments were conducted by following [3]. On LPBA40, we train our model on 30 subjects, generating 30×29 volume pairs, and test on the remaining 10 subjects. We follow [3] to merge 56 labels into 7 regions on the LPBA40, and center-crop the volumes into a size of $160 \times 192 \times 160$. On Mindboggle101, the data was divided into 42 subjects (with 1722 pairs) for training, and 20 subjects with 380 pairs for testing. All volumes were cropped to $160 \times 192 \times 160$. For evaluation metrics, we adopted the Dice score, which

measures the degree of overlap at the voxel level. The proposed Dual-PRNet was implemented in Pytorch and trained on 4 Titan Xp GPUs. Batch size was set to 4, due to the limitation of GPU memory. We adopt Adam optimization with a learning rate of $1e-4$.

3.1 Performance on Large Displacements




Fig. 2. **Top:** the generated deformation fields with increasing resolutions (left→right), and the last one is the final deformation field. **Bottom:** left→right: the moving image, the warped images by four deformation fields, and the fixed image.

Visualization. We first visualize the generated multi-resolution deformation fields in Fig. 2 (top). As can be found, the deformation field generated from a lower-resolution layer contains coarse and high-level context information, which is able to warp a volume at a larger scale. Conversely, the deformation field estimated from a higher-resolution layer can capture more detailed features. Fig. 2 (bottom) shows the warped images by four deformation fields presented. The warped images are refined gradually toward the fixed image, by aggregating more detailed structural information. We investigate the capability of our Dual-PRNet for handling large displacements, and compare our registration results against that of VoxelMorph in Fig. 3. Our Dual-PRNet can align the image more accurately than VoxelMorph, especially on the regions containing large displacements, as indicated in green or red regions.

On Large Spacing Displacements. We further evaluate the performance of Dual-PRNet on large spacing displacements. Experiments were conducted on LPBA40, by reducing the slices of the moving volumes from $160 \times 192 \times 160$ to $160 \times 24 \times 160$. During testing, the estimated final deformation field is applied to the labels of the moving volume using zero-order interpolation. With a significant reduction of slices from 192 to 24, our Dual-PRNet can still obtain a high average Dice score of 0.711, which even outperforms VoxelMorph [2] using the original non-reduction volumes, with an average Dice score of 0.683 achieved. This demonstrates the strong robustness of our model against large spacing displacements.




Fig. 3. Registration results on large displacements. From left to right: the moving image, the fixed image, results of VoxelMorph and Dual-PRNet.




Fig. 4. Dice scores of different methods on LPBA40 (7 regions).

3.2 Comparisons with the State-of-the-Art Approaches

We further compare our Dual-PRNet with a number of approaches: affine registration, SyN [10], and VoxelMorph [2], on LPBA40 and Mindboggle101 datasets. We implemented affine registration method and SyN by using ANTs [12]. For VoxelMorph, we used the codes and models provided by the original authors.

Results and Comparisons. On LPBA40 database, as shown in Fig. 4, our method obtains an average dice score of 0.778, and outperforms the others by a large margin, e.g., 0.731 of SyN and 0.683 of VoxelMorph. Our method achieves the best performance on all seven evaluated regions. The results on Mindboggle101 database are shown in Table 1, where our Dual-PRNet consistently outperforms the other three methods, and achieves best performance on all five regions. It reaches a high average Dice score of 0.631, compared to 0.525 of SyN and 0.511 of VoxelMorph.

Discussions. We provide ablation studies to further verify the efficiency of each technical component: the dual-stream design and pyramid registration module. Our dual-stream architecture, by estimating a single deformation field

Table 1. Dice scores of different methods on Mindboggle101.

Region	Frontal	Parietal	Occipital	Temporal	Cingulate	Average	Time(s)
Affine	0.455	0.406	0.354	0.469	0.450	0.427	--
SyN	0.558	0.496	0.446	0.578	0.549	0.525	--
VoxelMorph	0.532	0.459	0.480	0.585	0.499	0.511	0.451
Dual-PRNet	0.602	0.690	0.550	0.695	0.618	0.631	1.724

as VoxelMorph, can achieve an average Dice score of 0.767 on the LPBA40, and 0.582 on the Mindboggle101, which improved the single-stream counterpart by +8.4% and +7.1% respectively. By integrating the pyramid registration module, the results are further increased, with +1.1% and +4.9% further improvements on the LPBA40 and Mindboggle101 respectively. Notice that the pyramid registration module has a larger improvement on the Mindboggle101, due to the more complicated anatomical structure provided in the Mindboggle101, which often require more accurate labels to identify subtle difference, and coarse-to-fine refinements of the deformation fields can naturally make more contribution.

4 Conclusion

We have presented a new Dual-Stream Pyramid Registration Network for unsupervised 3D medical image registration. We designed a two-stream 3D encoder-decoder network to compute two convolutional feature pyramids, and then proposed a pyramid registration module which predicts multi-scale registration fields directly from the decoding feature pyramids, allowing for recurrently refining the registration fields and convolutional features. This results in a high-performance model that can better handle large deformations in both spatial and spacing domains. With these technical improvements, our method achieved impressive performance for brain MRI registration, and significantly outperformed recent approaches.

References

1. Jaderberg, M., Simonyan, K., Zisserman, A.: Spatial transformer networks. In NIPS. pp. 2017-2025 (2015)
2. Balakrishnan, G., Zhao, A., Sabuncu, M. R., Guttag, J., Dalca, A. V.: An unsupervised learning model for deformable medical image registration. In CVPR. pp. 9252-9260 (2018)
3. Kuang, D., Schmah, T.: FAIM-A convnet method for unsupervised 3D medical image registration. arXiv preprint arXiv:1811.09243 (2018)
4. 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. pp.739-746. Springer, Cham(2018)
5. Lewis, K. M., Balakrishnan, G., Rost, N. S., Guttag, J., Dalca, A. V.: Fast Learning-based Registration of Sparse Clinical Images. arXiv:1812.06932 (2018)

6. Ranjan, A., Black, M. J.: Optical flow estimation using a spatial pyramid network. In CVPR. pp. 4161-4170 (2017)
7. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In MICCAI. pp. 234-241. Springer, Cham (2015)
8. Chen, H., Dou, Q., Yu, L., Qin, J., Heng, P. A.: VoxResNet: Deep voxelwise residual networks for brain segmentation from 3D MR images. *NeuroImage*, 170, 446-455 (2018)
9. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In CVPR. pp. 770-778 (2016)
10. Shattuck, D. W., Mirza, M., Adisetiyo, V., Hojatkashani, C., Salamon, G., Narr, K. L., Toga, A. W.: Construction of a 3D probabilistic atlas of human cortical structures. *Neuroimage*, 39(3), 1064-1080 (2008)
11. Klein, A., Tourville, J.: 101 labeled brain images and a consistent human cortical labeling protocol. *Frontiers in neuroscience*, 6, 171 (2012)
12. Avants, B. B., Tustison, N. J., Song, G., Cook, P. A., Klein, A., and Gee, J. C.: A reproducible evaluation of ants similarity metric performance in brain image registration. *Neuroimage*, 54(3), pp. 20332044 (2011)