



UNIVERSITY OF LEEDS

This is a repository copy of *A Deep Discontinuity-Preserving Image Registration Network*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/176533/>

Version: Accepted Version

Proceedings Paper:

Chen, X orcid.org/0000-0003-4203-4578, Xia, Y, Ravikumar, N et al. (1 more author) (2021) A Deep Discontinuity-Preserving Image Registration Network. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part IV. Medical Image Computing and Computer Assisted Intervention – MICCAI 2021, 27 Sep - 01 Oct 2021, Strasbourg, France/Virtual. Springer, Cham , pp. 46-55. ISBN 978-3-030-87202-1

https://doi.org/10.1007/978-3-030-87202-1_5

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

A Deep Discontinuity-Preserving Image Registration Network

Xiang Chen¹, Yan Xia^{1,2}, Nishant Ravikumar^{1,2*}, and Alejandro F Frangi^{1,2,3,4*}

¹ Center for Computational Imaging and Simulation Technologies in Biomedicine, School of Computing, University of Leeds, Leeds, UK

² Biomedical Imaging Department, Leeds Institute for Cardiovascular and Metabolic Medicine, School of Medicine University of Leeds, Leeds, UK

³ Department of Cardiovascular Sciences, KU Leuven, Leuven, Belgium

⁴ Department of Electrical Engineering, KU Leuven, Leuven, Belgium

⁵ * Indicates joint last authorship.

Abstract. Image registration aims to establish spatial correspondence across pairs, or groups of images, and is a cornerstone of medical image computing and computer-assisted-interventions. Currently, most deep learning-based registration methods assume that the desired deformation fields are globally smooth and continuous, which is not always valid for real-world scenarios, especially in medical image registration (e.g. cardiac imaging and abdominal imaging). Such a global constraint can lead to artefacts and increased errors at discontinuous tissue interfaces. To tackle this issue, we propose a weakly-supervised Deep Discontinuity-preserving Image Registration network (DDIR), to obtain better registration performance and realistic deformation fields. We demonstrate that our method achieves significant improvements in registration accuracy and predicts more realistic deformations, in registration experiments on cardiac magnetic resonance (MR) images from UK Biobank Imaging Study (UKBB), than state-of-the-art approaches.

Keywords: Deep Learning · Image Registration · Cardiac Image Registration · Discontinuity-preserving Image Registration.

1 Introduction

Image registration is a fundamental component of several applications in medical imaging. Recent years have seen a shift from traditional iterative methods to deep learning (DL)-based registration approaches. Although training DL-based approaches is time-consuming, inference is rapid, involving just a single forward pass through the network. Consequently, DL-based approaches offer substantial acceleration for pair-/group-wise image registration relative to traditional approaches, achieving near-real-time performance in certain applications.

Most existing DL-based registration methods constrain deformation fields to be globally smooth and continuous, through various means [3, 4, 7]. However, this

assumption is often violated in medical image registration applications, as tissue boundaries are naturally discontinuous. This is especially pronounced in cardiac or abdominal imaging, which involve large deformations of multiple tissue-types, and organ motion/sliding at tissue boundaries. Variability in the physical properties of different tissue-types results in discontinuities at native tissue boundaries [6, 5]. Hence, enforcing deformation fields to be globally smooth can generate unrealistic deformations and lead increased errors near these boundaries.

Discontinuity-preserving image registration is an active area of research in the context of traditional registration methods [15, 13, 11, 6]. For example, Hua et al. [6] proposed a discontinuous registration approach that utilised enriched B-spline basis functions at control points near discontinuous tissue boundaries, achieving significant improvement in registration accuracy, relative to other existing discontinuity-preserving registration methods. In contrast, only one study thus far has proposed a discontinuous DL-based image registration framework. Ng et al. [10] proposed a custom discontinuity-preserving regulariser on the deformation fields (used with a typical unsupervised registration network), to preserve discontinuities, while ensuring local smoothness within specific regions. They formulated a regularisation term based on the unsigned area of the parallelogram spanned by two displacement vectors associated with moving image voxels. However, without additional boundary information for guidance, such a discontinuity regularisation term alone is insufficient to preserve strong discontinuities in deformation fields.

This paper assumes that the desired deformation fields are locally smooth, but discontinuities may exist between different regions/organs at tissue interfaces. Therefore, we generate distinct smooth deformation fields for different regions of interest and compose them to obtain the final registration field, used to warp the moving image. Such a locally-smooth and globally-discontinuous registration scheme is achieved using a novel Deep Discontinuity-preserving Image Registration network, or DDIR. The contributions of this paper are two-fold: (1) we designed a novel framework, DDIR for discontinuous DL-based image registration. This is the first study to incorporate discontinuity in DL network structure and training strategy, and not only in terms of a custom regularisation term in the loss function. (2) Our proposed DDIR achieves significant improvement in registration accuracy over state-of-the-art registration methods, and preserves key cardiac morphological indices post-registration, not afforded by the latter.

2 Method

Pair-wise image registration aims to establish spatial correspondence between the moving image \mathbf{I}_M and fixed image \mathbf{I}_F and is formulated as,

$$\phi(\mathbf{x}) = \mathbf{x} + u(\mathbf{x}), \tag{1}$$

where, \mathbf{x} represents voxels/pixels in the moving image \mathbf{I}_M , $u(\mathbf{x})$ denotes the displacement field, and $\phi(\circ)$ represents the deformation function.

To generate deformation fields that are locally smooth and discontinuous at the boundaries of different organs/regions, we propose to generate deformation fields for different sub-regions, and compose them to obtain the final deformation field. Sub-regions in the images to be registered must first be segmented either manually or automatically. With short-axis (SAX) cardiac cine-magnetic resonance (CMR) images, manual and automatic segmentation results for left ventricle blood pool (LVBP), left ventricle myocardium (LVM) and right ventricle (RV) are generally available in public data sets, large-scale imaging initiatives (e.g. UK Biobank) and from previous studies on automatic CMR segmentation [2]. As the focus of this paper is on SAX-CMR image registration, we explicitly model discontinuities along cardiac boundaries by splitting the images into four sub-regions, namely, LVBP, LVM, RV, and background. These sub-regions are subsequently used to train our DDIR approach and register CMR images in manner that preserves discontinuities at their boundaries.

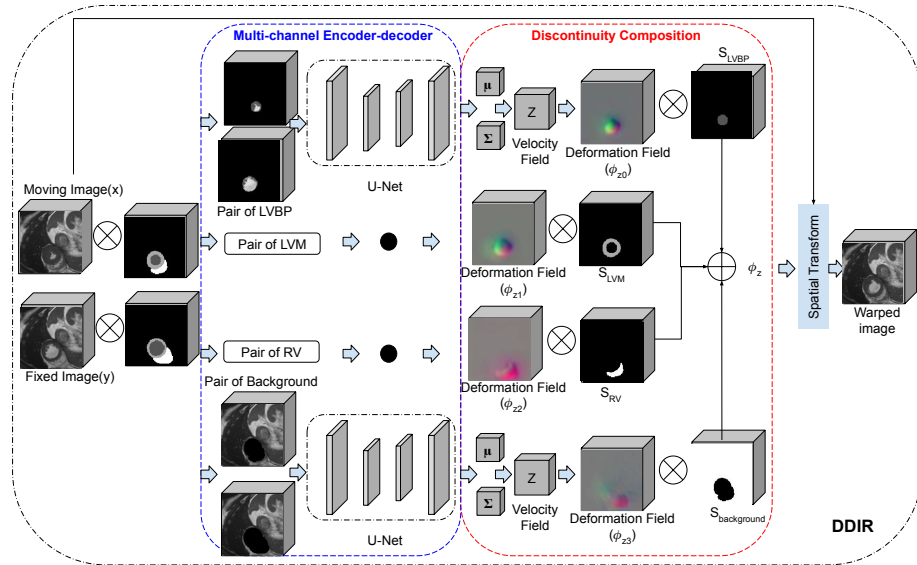


Fig. 1. Schema of DDIR. The registration network applies four different channels extracting features from pairs of LVBP, LVM, RV and background. Based on them, we obtain four sub-deformation fields for different regions. The final deformation field is obtained by composing these four deformation fields with corresponding segmentation. The cardiac MR images were reproduced by kind permission of UK Biobank ©.

Network Architecture. Most previous DL-based registration methods apply an encoder-decoder network (generally U-Net [12]) to extract feature maps from the concatenated input moving image and fixed image. However, as shown in Fig. 1, in DDIR the original moving image and fixed image (at $128 \times 128 \times 32$)

are divided into four image pairs, i.e. LVBP, LVM, RV and background, using segmentation masks for the corresponding regions. In each of these pairs, voxels in corresponding regions are preserved while the rest are set at zero. Each pair is concatenated and fed as input to a distinct U-Net block, which extracts region-specific feature maps. These four U-Nets have the same architecture, including four down-sampling layers and three corresponding up-sampling layers. Using this multi-channel encoder-decoder structure, we obtain four sets of feature maps ($64 \times 64 \times 16$) corresponding to different sub-regions. We use the same U-Net architecture (with identical hyper-parameters) in all DL-based registration approaches investigated in this study.

Discontinuity Composition. Using the region-specific feature maps learned by the U-Nets, we first predict four different smooth deformation fields (corresponding to each region) and then compose them to obtain the final deformation field, to preserve local smoothness and discontinuity at the interfaces. Similar to previous papers [4, 7], we assume the transformation function (denoted as ϕ_z) is parametrised by stationary velocity fields (SVF) ($z_i, i \in [0, 3]$), which are sampled from a multivariate Gaussian distribution. With the predicted feature map, we compute the mean μ_i and variance Σ_i of z_i (using two different convolution layers). Based on them, four SVFs (z_0, z_1, z_2, z_3) corresponding to different regions (LVBP, LVM, RV and background) are sampled. With the corresponding integration layer and up-sampling layer, we obtain four diffeomorphic deformation fields $\phi_{z_0}, \phi_{z_1}, \phi_{z_2}$ and ϕ_{z_3} . As before, we use region-specific segmentation masks to extract each region of interest from the obtained deformation fields (setting the remaining voxels to zero) and compose them to generate the final deformation field. Denoting the segmented regions of LVBP, LVM, RV and background as $S_{LVBP}, S_{LVM}, S_{RV}$ and $S_{background}$ respectively, the composition can be formulated as,

$$\phi_z = \phi_{z_0} \times S_{LVBP} + \phi_{z_1} \times S_{LVM} + \phi_{z_2} \times S_{RV} + \phi_{z_3} \times S_{background}. \quad (2)$$

Loss Function. The loss function includes two terms, a dissimilarity and a regularisation term. The former is the distance between the warped moving image and the fixed image, while, the latter constrains the estimated deformation fields to be locally smooth (i.e. within each region), to avoid unrealistic deformations. The dissimilarity loss in DDIR captures the dissimilarity on both images and segmentations. We use normalised cross-correlation (NCC) L_{NCC} to evaluate the similarity between the warped moving image and the fixed image. As the region-wise segmentation masks are available, we also compute the region-wise dice loss, denoted L_{Dice} as in [9].

To preserve discontinuity at the interfaces of the organs/regions while ensuring local smoothness, a global smoothness constraint is not enforced on the composed deformation field. The composition of different deformation fields preserves discontinuities at interfaces, therefore, we only need to guarantee the deformation field of each sub-region smooth. This is achieved by regularising each sub-deformation field. Following Voxelmorph-diff [4], we calculate the Kullback-Leibler (KL) divergence between the approximate posterior $q_\psi(z|\mathbf{I}_F; \mathbf{I}_M)$ and

the prior $p(z)$ ($p(z) = \mathcal{N}(z; 0, \Sigma_z)$) of each velocity field z , formulated as,

$$\begin{aligned} R &= KL(q_\psi(z|\mathbf{I}_F; \mathbf{I}_M)||p(z|\mathbf{I}_F; \mathbf{I}_M)), \\ L_R &= \frac{1}{4}(R_{LVBP} + R_{LVM} + R_{RV} + R_{background}), \end{aligned} \quad (3)$$

where R denotes the regularisation for each deformation field and L_R is the combined regularisation term. The $q_\psi(z|\mathbf{I}_F; \mathbf{I}_M) = \mathcal{N}(z; \mu_{\mathbf{z}|\mathbf{I}_F, \mathbf{I}_M}, \Sigma_{\mathbf{z}|\mathbf{I}_F, \mathbf{I}_M})$ is a multivariate normal, where, $\mu_{\mathbf{z}|\mathbf{I}_F, \mathbf{I}_M}$ and $\Sigma_{\mathbf{z}|\mathbf{I}_F, \mathbf{I}_M}$ are the mean and variance of the distribution, learned by convolution layers. The complete loss function used to train the network is, $L_{total} = \lambda_0 \times L_{NCC} + \lambda_1 \times L_{Dice} + \lambda_2 \times L_R$, where, λ_0 , λ_1 and λ_2 are used to weight the importance of each loss term.

3 Experiments and Results

Data and Implementation. The registration performance of the proposed approach is evaluated on SAX-CMR images (spatial resolution at $\sim 1.8 \times 1.8 \times 10mm^3$), available from UKBB. We chose images from 2,000 subjects at random, and used images at end-diastole (ED) and end-systole (ES) for intra-subject registration. Among these, 1,600 subjects’ data was chosen at random for training DDIR, equating to 3,200 image pairs (ED-to-ES or ES-to-ED registration). Image pairs from the remaining 400 subjects were used for testing. All CMR images were resampled to $1.50 \times 1.50 \times 3.15mm^3$ using bi-cubic interpolation, and cropped to a size of $128 \times 128 \times 32$ (with zero-padding for images with fewer than 32 slices). The region-wise segmentation masks for all CMR images were obtained automatically using the segmentation method proposed in [2]. DDIR was implemented using Python and Keras on a Tesla M60 GPU machine. The Adam optimiser was used for training, with a learning rate of $1e-4$. The batch size was set to 2, and the hyper-parameters λ_0 , λ_1 and λ_2 were set to 20, 200, 0.1 (determined empirically), respectively. The source code will be publicly available on the Github ⁶.

Quantitative Comparison and Analysis. To demonstrate the superiority of our approach, we compare DDIR with both traditional registration and DL-based registration methods. For the former, we choose Symmetric Normalisation (SyN) registration (3 resolution level, with 100 iterations in each sampling level) in ANTS [1], Demons (Fast Symmetric Forces Demons [14] with 800 iterations and standard deviations 1.0) in SimpleITK and B-spline registration (max iteration step is 2000, sampling 6000 random points per iteration) in SimpleElastix [8], for comparison. For the latter, DDIR is compared with Voxelmorph-diff [4]. As DDIR uses segmentation masks during training and inference, it is a weakly-supervised registration method. For fair comparison, we build three weakly-supervised versions of Voxelmorph - VM-Dice, VM(img+seg) and VM-Dice(img+seg). VM-Dice uses a Dice loss L_{Dice} term and binary cardiac segmentation masks for the fixed and moving images during training, but

⁶ <https://github.com/cistib/DDIR>

does not require the latter for inference. In VM(img+seg), we concatenate the fixed and moving images with their corresponding multi-class masks (i.e. distinct labels for each region) and use these to train the network. While, VM-Dice(img+seg) is a combination of the previous two methods. We did not compare with the DL-based discontinuity-preserving method proposed in [10], as there is no corresponding source code publicly available. This strategy to register different sub-regions and compose corresponding deformation fields is also applicable to the aforementioned networks. Hence, we also apply this strategy during inference, for trained Voxelmorph-diff and VM-Dice models (as they only require sub-images as input on the inference), for comparison with DDIR. These are denoted Voxelmorph-diff(compose) and VM-Dice(compose). These two approaches are different to DDIR as the composition of sub-deformation fields is not learned end-to-end during training (as in DDIR).

To demonstrate the advantage of incorporating discontinuity in the DL-based registration network, we also build a baseline for DDIR, DDIR(baseline), where the predicted feature maps from the four different channels are concatenated and used to compute a single diffeomorphic deformation field (instead of four sub-deformation fields, as in DDIR).

Qualitative Results. Registration results obtained using DDIR and the other methods investigated are assessed visually in Fig. 2. Here, the moving and fixed images are shown in the first column. The corresponding warped moving images, deformation fields, and Jacobian determinants (rows 1-3) obtained following registration using SyN, B-spline, Voxelmorph-diff, DDIR(baseline) and DDIR, are shown in columns 2-6. The warped moving images obtained by both traditional registration methods distinctly different to fixed image, although the B-spline result appears visually more similar than obtained by SyN. All warped moving images obtained using DL-based methods look more similar to the fixed image, than the former. The deformation fields and their corresponding Jacobian determinants estimated using each approach indicate that distinct boundaries for the left and right ventricle are retained using DDIR, not afforded by the rest.

Table 1. Quantitative comparison between DDIR and state-of-the-art methods using the DS of LVBP, LVM, RV and average Dice (denoted as Avg. DS) and HD. Statistically significant improvements in registration accuracy (DS and HD) are highlighted in bold. Besides, LVEDV and LVMM indices with no significant difference from the reference are also highlighted in bold.

Methods	LVBP DS (%)	LVM DS (%)	RV DS (%)	Avg. DS (%)	HD (mm)	LVEDV	LVMM
before Reg	57.68 ± 6.21	30.88 ± 8.68	55.13 ± 7.51	47.90 ± 6.33	12.91 ± 2.48	143.76 ± 32.13	83.67 ± 21.06
B-spline	74.44 ± 11.50	68.06 ± 7.20	61.76 ± 12.05	68.09 ± 8.76	13.72 ± 3.57	131.14 ± 40.64	81.11 ± 22.60
Demons	80.29 ± 10.00	69.96 ± 5.50	64.86 ± 9.67	71.70 ± 6.96	13.06 ± 3.12	138.00 ± 34.15	80.00 ± 21.25
SyN	70.92 ± 9.36	57.88 ± 10.59	60.30 ± 8.35	63.03 ± 8.29	12.98 ± 2.68	120.09 ± 41.83	83.12 ± 21.20
Voxelmorph-diff	81.73 ± 8.71	72.04 ± 4.65	65.73 ± 9.62	73.16 ± 6.26	12.96 ± 3.14	137.16 ± 32.59	78.65 ± 21.68
VM-Dice	82.28 ± 8.75	72.53 ± 4.59	66.30 ± 9.67	73.70 ± 6.28	13.00 ± 3.24	139.58 ± 32.79	78.98 ± 21.57
VM(img+seg)	82.54 ± 8.50	72.66 ± 4.80	66.69 ± 9.64	73.96 ± 6.28	12.68 ± 3.21	138.29 ± 33.00	80.83 ± 21.62
VM-Dice(img+seg)	81.97 ± 8.53	71.23 ± 4.79	70.20 ± 12.05	74.47 ± 6.79	11.28 ± 4.35	144.33 ± 32.93	80.17 ± 22.02
Voxelmorph-diff(compose)	78.82 ± 6.38	67.41 ± 8.80	75.10 ± 6.97	73.78 ± 6.10	11.74 ± 3.08	119.30 ± 38.71	91.39 ± 23.07
VM-Dice(compose)	79.59 ± 5.91	68.81 ± 7.81	77.93 ± 6.63	75.44 ± 5.36	11.14 ± 3.12	120.90 ± 38.14	94.89 ± 25.96
DDIR(baseline)	84.25 ± 8.63	75.02 ± 4.50	71.42 ± 10.32	76.90 ± 6.58	11.85 ± 3.38	141.73 ± 32.29	79.01 ± 21.40
DDIR	84.63 ± 8.07	75.27 ± 5.03	74.07 ± 8.73	77.99 ± 5.47	10.65 ± 3.51	141.84 ± 32.59	81.92 ± 21.86

Quantitative Results. To quantitatively evaluate the performance of our approach, we compare DDIR with previous methods using Dice score (DS) and the Hausdorff Distance (HD). DS is computed for LVBP, LVM and RV. These values and the average DS and HD across all regions are reported in Table 1. Besides, to demonstrate the clinical value of DDIR, we also compute two clinical indices, LV end-diastolic volume (LVEDV) and LV myocardial mass (LVMM). The former is computed using ED segmentations, while the latter, is computed using ED and ES segmentations, pre- and post-registration. Pre-registration, LVEDV and LVMM are computed based on the moving and fixed segmentations (used as reference values). Post-registration, we compute them based on the warped moving segmentation. Therefore, as we perform both ED-to-ES and ES-to-ED registration for each subject, the LVMM values reported in Table 1 represent the average computed at both ED and ES, across all subjects. Thus the closer LVEDV and LVMM (post-registration) are to the reference values, the better the registration performance.

DL-based approaches outperform traditional registration methods in terms of both DS and HD. The weakly-supervised variants of Voxelmorph-diff provide improvements over Voxelmorph-diff, consistent with previous research[4]. Using segmentation masks as additional input channels to the network (VM(img+seg)) yields better results than using them just to compute the loss and drive gradient




Fig. 2. Visual comparison of deformation fields estimated using DDIR and state-of-the-art methods. Left column: Moving and fixed images; Right column: corresponding warped moving image (first row), deformation fields (second row) and Jacobian Determinant (last row). Colours in the Jacobian determinant images, from blue to red represent the intensity from low to high. The cardiac MR images were reproduced by kind permission of UK Biobank ©.

updates (VM-Dice) (73.96% vs 73.70%). However, conversely the former requires segmentation masks during inference, while the latter do not. The combination of these two strategies (VM-Dice(img+seg)) further improves registration performance ($\sim 0.5\%$ in terms of average DS). Composing sub-deformation fields also improves registration accuracy of the trained networks, with Voxelmorph-diff (compose) achieving 0.6% higher average DS than Voxelmorph-diff (73.78% vs 73.16%), and VM-Dice (compose) achieving $\sim 1.7\%$ higher average DS than VM-Dice (75.44% vs 73.70%). We found that the DDIR(baseline) achieves $\sim 1\%$ higher average DS than VM-Dice(img+seg) (76.90% vs 75.93%), which highlights the advantage of using a multi-channel encoder-decoder network. Compared with DDIR, we found that incorporating discontinuity further improves the average DS (77.99% vs 76.90%). Correspondingly, DDIR also obtains the best performance in terms of the DS for LVBP, LVM and HD, while its RV DS is lower than VM-Dice(compose). We evaluated the statistical significance of these results using paired t-tests and found that DDIR significantly outperforms Voxelmorph-diff, VM-Dice, VM(img+seg) and VM-Dice(img+seg) on all DS and HD metrics (P-value <0.05). DDIR also significantly outperforms DDIR(baseline) in terms of average DS, RV DS and HD. Each sub-deformation field generated by DDIR are smooth (without foldings). After composing, the discontinuity only exists at the interface of different sub-regions, which demonstrates that DDIR can generate locally-smooth but globally-discontinuous deformation fields.

The clinical indices, LVEDV and LVMM, show no significant differences (P-value >0.05) post-registration using DDIR to the reference values, not afforded by other approaches. This demonstrates the superiority and clinical value of our method. To analyse the discontinuity on the deformation fields, we visualise the deformation fields generated using DDIR and DDIR (baseline) (presented in the supplementary material), where the discontinuity is observed for the former along the LV and RV boundaries. To further demonstrate the robustness and generalisability of our approach, we apply the models trained on UKBB data, to the publicly available Automatic Cardiac Diagnosis Challenge (ACDC) data set. The qualitative and quantitative results are included in the supplementary material for brevity. As cardiac motion in ACDC images is not as pronounced as in UKBB (in some cases, the images in ED are very similar to ES), only marginal differences in registration performance are observed between DDIR and the other composition-based methods in terms of DS and HD. However, as before, DDIR outperforms Voxelmorph-diff and traditional state-of-the-art methods. Additionally, the clinical indices quantified (LVEDV, LVMM) post registration using DDIR show no significant differences to the reference, not afforded by any of the other methods investigated. This demonstrates the potential for applying DDIR in real clinical scenarios.

4 Conclusion

We proposed a novel weakly-supervised discontinuity-preserving registration network, DDIR, which significantly outperformed the state-of-the-art, in intra-

patient CMR registration. DDIR preserves LV clinical indices post-registration, not afforded by the other approaches. This makes it compelling as a tool for use in clinical applications as it ensures that common diagnostic biomarkers for the LV are preserved post-registration.

Acknowledgements

This research was conducted using the UKBB resource under access application 11350 and was supported by the Royal Academy of Engineering under the RAEng Chair in Emerging Technologies (CiET1919/19) scheme and EPSRC TUSCA (EP/V04799X/1).

References

1. 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)
2. Bai, W., Sinclair, M., Tarroni, G., Oktay, O., Rajchl, M., Vaillant, G., Lee, A.M., Aung, N., Lukaschuk, E., Sanghvi, M.M., et al.: Automated Cardiovascular Magnetic Resonance Image Analysis with Fully Convolutional Networks. *Journal of Cardiovascular Magnetic Resonance* **20**(1), 65 (2018)
3. 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)
4. Dalca, A.V., Balakrishnan, G., Guttag, J., Sabuncu, M.R.: Unsupervised Learning of Probabilistic Diffeomorphic Registration for Images and Surfaces. *Medical Image Analysis* **57**, 226–236 (2019)
5. Hua, R.: Non-rigid Medical Image Registration with Extended Free Form Deformations: Modelling General Tissue Transitions. Ph.D. thesis, University of Sheffield (2016)
6. Hua, R., Pozo, J.M., Taylor, Z.A., Frangi, A.F.: Multiresolution eXtended Free-Form Deformations (XFFD) for Non-rigid Registration with Discontinuous Transforms. *Medical Image Analysis* **36**, 113–122 (2017)
7. Krebs, J., Delingette, H., Mailhé, B., Ayache, N., Mansi, T.: Learning A Probabilistic Model for Diffeomorphic Registration. *IEEE Transactions on Medical Imaging* **38**(9), 2165–2176 (2019)
8. Marstal, K., Berendsen, F., Staring, M., Klein, S.: SimpleElastix: A User-friendly, Multi-lingual Library for Medical Image Registration. In: *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*. pp. 134–142 (2016)
9. Milletari, F., Navab, N., Ahmadi, S.A.: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: *2016 fourth international conference on 3D vision (3DV)*. pp. 565–571. IEEE (2016)
10. Ng, E., Ebrahimi, M.: An Unsupervised Learning Approach to Discontinuity-Preserving Image Registration. In: *International Workshop on Biomedical Image Registration*. pp. 153–162. Springer (2020)

11. Pace, D.F., Aylward, S.R., Niethammer, M.: A Locally Adaptive Regularization based on Anisotropic Diffusion for Deformable Image Registration of Sliding Organs. *IEEE transactions on Medical Imaging* **32**(11), 2114–2126 (2013)
12. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional Networks for Biomedical Image Segmentation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 234–241. Springer (2015)
13. Schmidt-Richberg, A., Werner, R., Handels, H., Ehrhardt, J.: Estimation of Slipping Organ Motion by Registration with Direction-dependent Regularization. *Medical Image Analysis* **16**(1), 150–159 (2012)
14. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N., et al.: Diffeomorphic Demons Using ITK's Finite Difference Solver Hierarchy. *The Insight Journal* **1** (2007)
15. Wu, Z., Rietzel, E., Boldea, V., Sarrut, D., Sharp, G.C.: Evaluation of Deformable Registration of Patient Lung 4DCT with Subanatomical Region Segmentations. *Medical Physics* **35**(2), 775–781 (2008)