

TEDS-Net: Enforcing Diffeomorphisms in Spatial Transformers to Guarantee Topology Preservation in Segmentations

Madeleine K. Wyburd¹, Nicola K. Dinsdale², Ana I.L. Namburete^{†1} and Mark Jenkinson^{†2,3,4}

¹ Institute of Biomedical Engineering, Department of Engineering Science, University of Oxford, Oxford, UK

² Wellcome Centre for Integrative Neuroimaging, FMRIB, University of Oxford, Oxford, UK

³ Australian Institute for Machine Learning (AIML), Department of Computer Science, University of Adelaide, Adelaide, Australia

⁴ South Australian Health and Medical Research Institute (SAHMRI), North Terrace, Adelaide, Australia

`madeleine.wyburd@hertford.ox.ac.uk`

Abstract. Accurate topology is key when performing meaningful anatomical segmentations, however, it is often overlooked in traditional deep learning methods. In this work we propose TEDS-Net: a novel segmentation method that guarantees accurate topology. Our method is built upon a continuous diffeomorphic framework, which enforces topology preservation. However, in practice, diffeomorphic fields are represented using a finite number of parameters and sampled using methods such as linear interpolation, violating the theoretical guarantees. We therefore introduce additional modifications to more strictly enforce it. Our network learns how to warp a binary prior, with the desired topological characteristics, to complete the segmentation task. We tested our method on myocardium segmentation from an open-source 2D heart dataset. TEDS-Net preserved topology in 100% of the cases, compared to 90% from the U-Net, without sacrificing on Hausdorff Distance or Dice performance. Code will be made available at: www.github.com/mwyburd/TEDS-Net.

Keywords: Diffeomorphic · Topology · Segmentation · Spatial Transformers

1 Introduction

Anatomical segmentations with incorrect topology can be clinically problematic and highly impractical. For example, when segmenting the heart’s myocardium, topological mistakes may present as isolated segmented regions, small holes within the wall or a disconnected perimeter. These inaccuracies can all lead to incorrect measurements of circumference and wall thickness, which are often required when diagnosing heart conditions such as hypertrophic cardiomyopathy [11]. Post-processing morphological operations, such as binary closing, can

be used to correct some small gaps and disconnected regions. However, they are unaware of the structure’s overall topology and need to be customised accordingly, therefore, an unified method is preferred.

Semantic segmentation of medical images is regularly performed using deep convolutional neural networks (CNNs), such as the U-Net [15]. One alternative segmentation method uses a combination of CNNs and spatial transformers to learn the spatial warping required to transform a set of prior shapes into the desired class labels [10,18,17,8,19]. Such methods have outperformed conventional encoder-decoder and state-of-the-art architectures, however, they have no topological guarantees.

Moreover, both these methods are commonly trained and evaluated using loss functions such as Dice or binary cross-entropy, which evaluate each pixel individually and not the higher level structure, often resulting in large topological errors. Recent work has recognised the importance of accurate whole-structure segmentation, with the development of loss functions that encourage the preservation of topology [5,14,9]. Although these techniques report an improved topology accuracy compared to standalone pixel-wise loss functions, they only encourage topology preservation as opposed to enforcing it, with the best performing method achieving 96.7% topology accuracy in myocardium segmentation [5].

In theory, topology will always be preserved by deforming a prior, with the correct topological characteristic, with a continuous *diffeomorphic field* [2]. Diffeomorphic fields are continuous deformation fields that result in a one-to-one mapping. Their derivatives are invertible, resulting in positive Jacobian determinants, giving an unambiguous mapping between coordinates and therefore preserving topology [2].

Resampling one diffeomorphic field by another, also known as a *composition*, always results in a new diffeomorphic field [2]. In VoxelMorph’s methods [6,7], this elegant property has been utilised for brain registration, by initialising with a small diffeomorphic field and amplifying it through a series of novel integration layers, based on the squaring and scaling approach [1,13], to generate a topology-preserving large-deformation field. However, when diffeomorphisms are applied in discrete settings, such as images stored as discrete voxels, their topological guarantees can begin to break down. This is often indicated by the emergence of non-positive Jacobian determinants in the deformation fields, which correspond to folding voxels in the warped space [7]. In brain registration a small fraction of violations may be manageable, conversely, for segmentation tasks these can lead to topological errors in the structure of interest, defeating the purpose. Here, we adapt VoxelMorph’s methods to be suitable for segmentation tasks by introducing a novel topology-preserving activation function and additional smoothing terms to counteract the negative effects of discrete sampling. These modifications enforce true diffeomorphisms, shown to have only positive Jacobian determinants, which results in 100% topology preservation.

In this work, we present the novel Topology Enforcing Diffeomorphic Segmentation Network (TEDS-Net), which to the best of our knowledge is the first deep learning technique to achieve 100% topology accuracy, and to combine spatial

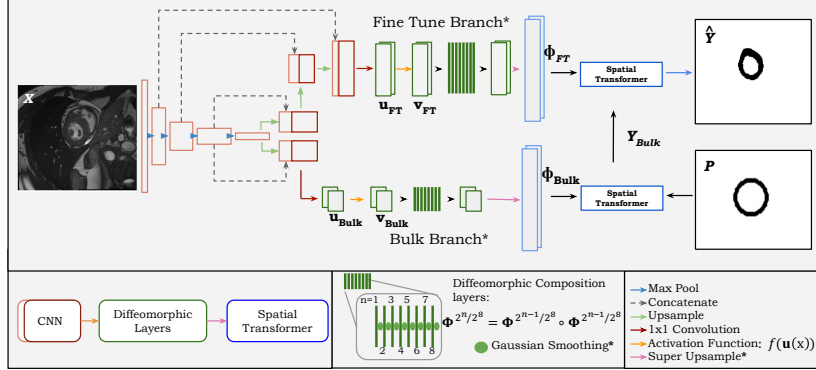


Fig. 1. Schematic of TEDS-Net architecture. A CNN learns two initial fields, \mathbf{u} , at different resolutions, from an input image \mathbf{X} . The fields are enforced to be diffeomorphic using an activation function, $\mathbf{v} = f(\mathbf{u})$, amplified through composition layers and “super” upsampled to 2x the resolution of the input. The bulk displacement, Φ_{Bulk} , samples a binary prior, \mathbf{P} , generating \mathbf{Y}_{Bulk} , which is then sampled by the fine tuning field, Φ_{FT} . The asterisks show the elements removed during ablation studies.

transformer networks (STN) and diffeomorphic displacement fields to complete a segmentation as the primary task. Our method uses CNNs to learn a warp field that maps a binary prior on to the input image, with the warp field being diffeomorphic by construction. We show that in the discrete setting, continuous properties of diffeomorphisms are no longer guaranteed, therefore, we introduce model components that enforce topology preservation in all cases.

2 Method

The aim of the proposed TEDS-Net is to predict a segmentation label, $\hat{\mathbf{Y}}$, using learnt diffeomorphic transformations, Φ , applied to a binary prior image, \mathbf{P} . The network is comprised of three main parts, shown in the bottom left corner of Fig 1. Firstly, an encoder-decoder style network is used to learn the initial velocity fields. These fields are then enforced to be diffeomorphic and amplified using the squaring and scaling approach [13,1]. Finally, the diffeomorphic fields are used to sample a prior binary shape to complete the segmentation task.

CNN. An encoder-decoder network was used to extract the relevant features from an input image, $\mathbf{X} \in \mathbb{R}^{h \times w}$ with dimensions $[h \times w]$, in order to predict two initial velocity fields: a *bulk field* and a *fine-tuning field*. We used an architecture similar to the U-Net but with two decoder branches [15,16]. Each branch consisted of a series of convolutions followed by an instance normalisation and ReLU activation function and then repeated, referred to as a convolutional block.

Branching off at the bottleneck were two decoder streams, made up of convolutional blocks and a final 1x1 convolution, used to generate a two-dimensional

field, \mathbf{u} . The first stream, with one upsampling convolutional block, predicts a low resolution bulk-velocity, $\mathbf{u}_{\text{Bulk}} \in \mathbb{R}^{2 \times (h \times w)/8}$, which focuses on warping the prior shape into the correct region. The second, with three upsampling blocks, predicts a higher resolution fine-tuning field, used to fine-tune the warped shape, $\mathbf{u}_{\text{FT}} \in \mathbb{R}^{2 \times (h \times w)/2}$.

Diffeomorphic Layers. To enforce that the initial fields, \mathbf{u} , were diffeomorphic, and therefore suitable for the scaling and squaring approach, a customised tanh activation function, $f(\mathbf{u})$, was applied to the output of each decoder branch to enforce that the displacements were between -0.5 voxel and $+0.5$ voxel and therefore topology preserving:

$$\mathbf{v}(\mathbf{x}) = f(\mathbf{u}(\mathbf{x})) = 0.5 \left(\frac{e^{\mathbf{u}(\mathbf{x})} - e^{-\mathbf{u}(\mathbf{x})}}{e^{\mathbf{u}(\mathbf{x})} + e^{-\mathbf{u}(\mathbf{x})}} \right). \quad (1)$$

A series of composition layers were then used to amplify the initial diffeomorphic fields, \mathbf{v} , adapted from the VoxelMorph’s implementation [6,7].

Gaussian Smoothing: As the diffeomorphic fields are represented using a finite number of parameters and sampled using linear interpolation, the theoretical guarantees are violated, risking key properties such as topological preservation. Moreover, these inaccuracies and imperfections can be amplified with the number of compositions performed [2]. To reduce these violations, we introduced Gaussian smoothing between each integration layer. For this work we used a 5 by 5 kernel and $\sigma = 2$ and the impact of this will be investigated in future work.

Super Upsample: To further smooth the fields and minimise any abnormalities, the amplified diffeomorphic flow fields were both “super upsampled” to double the resolution of the input image, \mathbf{X} , using bilinear interpolation.

Spatial Transformers. The diffeomorphic field, $\Phi_{\text{Bulk}} \in \mathbb{R}^{2 \times 2h \times 2w}$, generated from the first decoder branch sampled the prior shape, $\mathbf{P} \in \mathbb{R}^{h \times w}$, using bilinear interpolation: $\mathbf{Y}_{\text{Bulk}} = \Phi_{\text{Bulk}}(\mathbf{P})$, as shown in Fig 2. The resulting warped image was then sampled by the fine-tuning field generated by the second decoder branch, $\hat{\mathbf{Y}} = \Phi_{\text{FT}}(\mathbf{Y}_{\text{Bulk}})$, where $\Phi_{\text{FT}} \in \mathbb{R}^{2 \times 2h \times 2w}$. The prediction, $\hat{\mathbf{Y}}$, was then downsampled to the same resolution as the label, $\mathbf{Y} \in \mathbb{R}^{h \times w}$, using Max Pooling.

Field Regularisation: The network outputs the final label prediction, $\hat{\mathbf{Y}}$, and the two diffeomorphic transformation fields: Φ_{Bulk} , Φ_{FT} , which were all used for the training loss. A combination of Dice Loss ($\mathcal{L}_{\text{Dice}}$) and a field regularisation ($\mathcal{L}_{\text{Grad}}$) function were used for training:

$$\mathcal{L} = \mathcal{L}_{\text{Dice}}(\mathbf{Y}, \hat{\mathbf{Y}}) + \beta \mathcal{L}_{\text{Grad}}(\Phi_{\text{FT}}) + \beta \mathcal{L}_{\text{Grad}}(\Phi_{\text{Bulk}}), \quad (2)$$

where $\mathcal{L}_{\text{Grad}} = \sum_{i,j=1}^{2h,2w} \|\nabla \Phi(i,j)\|^2$, which encourages smooth flow fields by penalising large spatial gradients between neighbouring voxels, adapted from [3].




Fig. 2. A prior is sampled with a displacement field at double the resolution, $\mathbf{Y}_{\text{Bulk}} = \Phi_{\text{Bulk}}(\mathbf{P})$, before the second super upsampled displacement is applied to the shifted image $\hat{\mathbf{Y}} = \Phi_{\text{FT}}(\mathbf{Y}_{\text{Bulk}})$. The displacement panels show the shift in both directions.

A weighting parameter, β , was used to balance the contributions between $\mathcal{L}_{\text{Grad}}$ and $\mathcal{L}_{\text{Dice}}$, in this work we used a constant β of 10,000, in order to give all the loss function components similar magnitudes, as measured empirically.

2.1 Experimental Setup

TEDS-Net was used to segment the myocardium from an open source dataset of MRI heart slices in the short axis⁵ [4]. This dataset consisted of 2 scans from 100 patients, across five different pathologies. Five myocardium-containing slices were taken from each scan, cropped to 144 by 208 pixels and augmented using rotations, shifts, zooming and a combination of the three, resulting in 8,000 images split 75%, 15% and 10% for training, validation and testing, respectively. It was assured that slices from the same patient were assigned to the same subset.

The topology of the myocardium is equivalent to a hollow circle, which we used as the prior shape, \mathbf{P} , shown in Fig 2. It should be noted that the patients' scans were not aligned and therefore the location of the myocardium was inconsistent. Additionally, different pathologies have varying myocardium thickness. Despite this, the same arbitrary prior shape was used throughout, illustrating the robustness of the method.

Two main experiments were completed and evaluated on 100 unseen heart slices without augmentations. Firstly, ablation studies were performed to show the contribution of each additional element used in TEDS-Net to tackle the limitations of diffeomorphic sampling, shown with the asterisks in Fig 1, using a prior shape of radius 30 and 8 integration layers, which were empirically set. As a comparison, we used a U-Net [15] and VoxelMorph diffeomorphic network [6,7], using the shape prior \mathbf{P} as the second input channel. Secondly, we investigated the effect of the number of integration layers and radius of the binary prior.

The networks were all trained end-to-end on a NVIDIA GeForce GTX 1080 GPU and implemented in Pytorch using Python 3.6. Training was performed over 200 epochs, using the Adam optimiser with a learning rate of 0.0001 and a batch size of 5. TEDS-Net and the U-Net were both made up of 5 layers, with 12 initial feature maps, and a 20% dropout applied to each map.

⁵ The ACDC database: www.creatis.insa-lyon.fr/Challenge/acdc

Table 1. Comparison of myocardium segmentation performed by the baselines and TEDS-Net, accompanied by ablation studies (A1-6). The mean and standard deviation are given for Dice, Hausdorff Distance (HD) and percentage of non-positive Jacobian determinants from each field. The best result from each measure is shown in bold.

Network	Dice	HD	$\% J_{\Phi_{\text{Bulk}}} \leq 0$	$\% J_{\Phi_{\text{FT}}} \leq 0$	Incorrect Topology
U-Net (baseline)	0.87 ± 0.16	5.25 ± 13.11	N/A	N/A	10/100
VoxelMorph (baseline)	0.79 ± 0.18	3.0 ± 2.4	N/A	$6.7e^{-05} \pm 0.7e^{-03}$	6/100
(A1) TEDS-Net: No Gaussian, $\mathcal{L}_{\text{Grad}}$ or Super Upsample	0.82 ± 0.19	4.26 ± 5.76	56.32 ± 1.97	43.8 ± 3.46	17/100
(A2) TEDS-Net: No Gaussian Smoothing	0.86 ± 0.13	3.73 ± 8.34	67.10 ± 7.07	35.81 ± 4.41	3/100
(A3) TEDS-Net: No $\mathcal{L}_{\text{Grad}}$	0.85 ± 0.14	2.77 ± 2.01	5.94 ± 2.57	22.15 ± 3.67	35/100
(A4) TEDS-Net: No Super Upsample	0.85 ± 0.11	2.89 ± 1.61	0.0 ± 0.0	0.0 ± 0.0	0/100
(A5) TEDS-Net: Only Bulk Branch	0.77 ± 0.15	3.86 ± 3.50	0.03 ± 0.17	N/A	0/100
(A6) TEDS-Net: Only Fine-Tune Branch	0.85 ± 0.14	3.01 ± 2.16	N/A	0.0 ± 0.0	0/100
TEDS-Net (ours)	0.86 ± 0.12	2.76 ± 2.13	0.0 ± 0.0	0.0 ± 0.0	0/100

3 Results and Discussion

To evaluate segmentation performance, Dice, Hausdorff Distance (HD), and a count of incorrect topologies were used, as seen in Table 1. A predicted segmentation’s topology was classed as incorrect if their Betti numbers [12] varied from the known topological properties of the myocardium. To assess the discrete fields, the Jacobian determinants were also measured, as perfect diffeomorphic fields always have strictly positive determinants.

Large topological errors were found when segmenting the myocardium with the U-Net, as seen from Table 1. These errors were mainly expressed as large gaps separating the myocardium, shown in Fig 3, making clinical measures of circumference extremely challenging. VoxelMorph was designed to maximise registration performance using diffeomorphisms, however, when applied to segmentation a small fraction of topology errors are observed, shown in Table 1: VoxelMorph. Unlike the U-Net predictions, these topology violations are seen at the boundaries of the myocardium, forming holes or disconnected regions when defined using 4-pixel connectivity, as done here, but not when using 8-pixel connectivity. Conversely, our TEDS-Net enforced diffeomorphisms, preserving topology for all cases, which to the best of our knowledge, is the first deep learning segmentation technique to achieve 100% topological accuracy. Therefore, TEDS-Net additional modifications are required to prioritise topology preservation.

Paired t-tests were computed between TEDS-Net and U-Net and found that TEDS-Net significantly outperformed the U-Net in Hausdorff Distance, with a p-value of 0.01. Although the Dice accuracies were competitive, the U-Net performed significantly better (p=0.04). Figure 3c-d shows particularly challenging examples, where the U-Net fails to return the correct topology. However, although TEDS-Net returns labels with accurate topology, the myocardial wall appears too thick in some parts. This is likely due to the smoothing modification, whose parameter will be further investigated in future work.




Fig. 3. Examples of myocardium manual annotations (green), compared to U-Net and TEDS-Net segmentations (pink) from the ACDC dataset [4].

Ablation studies were performed to show the effect of the additional smoothing modifications used to encourage perfect diffeomorphisms, as shown in Table 1 (A1-6). The continuous guarantees of diffeomorphic sampling, such as topology preservation, break down when applied in the discrete setting, shown by the emergence of non-positive Jacobian determinants that can then lead to topological errors. There are two sources of such violations in TEDS-Net: the composition layers and the image sampling.

To limit the numerical inaccuracies brought about by the discrete composition of two diffeomorphic fields, we introduced Gaussian smoothing between each layer. Without this addition, a large fraction of the Jacobian determinants no longer remained positive, corresponding to a number of topological defects, as shown in Table 1 (A2).

Sampling a discrete binary image with a finite warp field often results in disconnects, due to interpolations and the use of thresholds on the resultant image. To reduce this effect, we regularised the smoothness of the final deformation fields with $\mathcal{L}_{\text{Grad}}$. Without this term, 35 out of 100 images were found to have topological defects and therefore, it played a vital role in mimicking perfect diffeomorphisms, as shown in Table 1 (A3).

Removing either the bulk or fine-tuning sampling branch was found to reduce the performance of both Dice and HD, as shown Table 1 (A5-6). As the images are unaligned, with varying sizes of myocardium, shown in Fig 3, both branches are required to first align the prior before fine-tuning the warped shape. TEDS-Net performs the worst without the fine-tuning branch, which is likely due to the lack of flexibility in the bulk transformation, as the deformation fields are generated at a much lower resolution before being upsampled.




Fig. 4. The effect that the number of integration layers (a) and the radius of the binary prior (b) had on segmentation performance and the diffeomorphic nature of the generated fields. Due to the image dimensions, 50 was the maximum radius used.

Number of Integration Layers: Including Gaussian smoothing between the integration layers played a key role in enforcing topology preservation, as shown in Table 1 (A2). Although theoretically accuracy should increase with the number of composition layers, this requires performing more compositions that can each bring about small violations, due to the discrete nature of the fields [2]. However, limiting the number of composition layers, limits the size of deformations. To investigate this further, we varied the number of integration layers used in TEDS-Net and measured the segmentation performance and Jacobian determinants of the resulting fields, shown by Fig 4a. When using 12 or fewer integration layers, the segmentation performance is stable whilst the Jacobian determinants all remain positive and all topology is preserved. However, when the number of layers is increased beyond 12, non-positive Jacobian determinants emerge. Therefore, whilst Gaussian smoothing has been shown to be essential in the integration layers, there is a limit to the number of compositions that can be used whilst enforcing diffeomorphic sampling.

Prior Radius: In the test set, the myocardium’s radius ranged between approximately 10 to 60 voxels, so we investigated the impact of varying the radius on the final segmentation performance, as shown in Fig 4b. The segmentation performance remained consistent as the radius was increased between 20 and 50 voxels. However, when the radius was set to 10 voxels, a small percentage of the Jacobian determinants of the fine-tuning fields were non-positive, corresponding to 2/100 incorrect topologies. This is likely due to the decreased number of voxels in the prior representing the myocardium and the area enclosed by it, in combination with the additional smoothing modification that restrict the flexibility of the diffeomorphic fields.

4 Conclusion

We have shown that TEDS-Net is able to achieve highly accurate myocardium segmentations whilst ensuring topology preservation. We introduced additional diffeomorphic-encouraging modifications, which were found to play a crucial role in enforcing an one-to-one mappings in the generated discrete fields. Our method successfully segmented the myocardium in unaligned MRI heart slices, with different pathologies that had different thicknesses and circumferences, using the same general prior shape for all. This flexible, easy to train method has the potential to have a high impact in future clinical segmentation work.

Acknowledgements:

MW and ND are supported by the Engineering and Physical Sciences Research Council (EPSRC) and Medical Research Council (MRC) [grant number EP/L016052/1]. MJ is supported by the National Institute for Health Research (NIHR) Oxford Biomedical Research Centre (BRC), and this research was funded by the Wellcome Trust [215573/Z/19/Z]. The Wellcome Centre for Integrative Neuroimaging is supported by core funding from the Wellcome Trust [203139/Z/16/Z]. AN is grateful for support from the UK Royal Academy of Engineering under the Engineering for Development Research Fellowships scheme.

References

1. Arsigny, V., Commowick, O., Pennec, X., Ayache, N.: A log-euclidean polyaffine framework for locally rigid or affine registration. In: International Workshop on Biomedical Image Registration. pp. 120–127. Springer (2006)
2. Ashburner, J.: A fast diffeomorphic image registration algorithm. *Neuroimage* **38**(1), 95–113 (2007)
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. Bernard, O., Lalande, A., Zotti, C., Cervenansky, F., Yang, X., Heng, P.A., Cetin, I., Lekadir, K., Camara, O., Ballester, M.A.G., et al.: Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis: is the problem solved? *IEEE transactions on medical imaging* **37**(11), 2514–2525 (2018)
5. Clough, J., Byrne, N., Oksuz, I., Zimmer, V.A., Schnabel, J.A., King, A.: A topological loss function for deep-learning based image segmentation using persistent homology. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020)
6. 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)
7. 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)

8. Dinsdale, N.K., Jenkinson, M., Namburete, A.I.: Spatial warping network for 3d segmentation of the hippocampus in mr images. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 284–291. Springer (2019)
9. Hu, X., Fuxin, L., Samaras, D., Chen, C.: Topology-preserving deep image segmentation. arXiv preprint arXiv:1906.05404 (2019)
10. Jaderberg, M., Simonyan, K., Zisserman, A., et al.: Spatial transformer networks. *Advances in neural information processing systems* **28**, 2017–2025 (2015)
11. Marian, A.J., Braunwald, E.: Hypertrophic cardiomyopathy: genetics, pathogenesis, clinical manifestations, diagnosis, and therapy. *Circulation research* **121**(7), 749–770 (2017)
12. Millson, J.J.: On the first betti number of a constant negatively curved manifold. *Annals of Mathematics* pp. 235–247 (1976)
13. Moler, C., Van Loan, C.: Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later. *SIAM review* **45**(1), 3–49 (2003)
14. Mosinska, A., Marquez-Neila, P., Koziński, M., Fua, P.: Beyond the pixel-wise loss for topology-aware delineation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3136–3145 (2018)
15. 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)
16. Stergios, C., Mihir, S., Maria, V., Guillaume, C., Marie-Pierre, R., Stavroula, M., Nikos, P.: Linear and deformable image registration with 3d convolutional neural networks. In: Image Analysis for Moving Organ, Breast, and Thoracic Images, pp. 13–22. Springer (2018)
17. Vigneault, D.M., Xie, W., Ho, C.Y., Bluemke, D.A., Noble, J.A.: ω -net (omega-net): fully automatic, multi-view cardiac mr detection, orientation, and segmentation with deep neural networks. *Medical image analysis* **48**, 95–106 (2018)
18. Wickramasinghe, U., Knott, G., Fua, P.: Probabilistic atlases to enforce topological constraints. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 218–226. Springer (2019)
19. Zeng, Q., Karimi, D., Pang, E.H., Mohammed, S., Schneider, C., Honarvar, M., Salcudean, S.E.: Liver segmentation in magnetic resonance imaging via mean shape fitting with fully convolutional neural networks. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 246–254. Springer (2019)