



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

# Generating Anatomically Accurate Heart Structures via Neural Implicit Fields

Jiancheng Yang<sup>1,\*</sup>[0000-0003-4455-7145], Ekaterina Sedykh<sup>1</sup>, Jason Ken Adhinarta<sup>2,\*\*</sup>, Hieu Le<sup>1</sup>, and Pascal Fua<sup>1</sup>[0000-0002-6702-9970]

<sup>1</sup> Swiss Federal Institute of Technology Lausanne (EPFL), Switzerland  
<sup>2</sup> Boston College, MA, USA

**Abstract.** Implicit functions have significantly advanced shape modeling in diverse fields. Yet, their application within medical imaging often overlooks the intricate interrelations among various anatomical structures, a consideration crucial for accurately modeling complex multi-part structures like the heart. This study presents ImHeart, a latent variable model specifically designed to model complex heart structures. Leveraging the power of learnable templates, ImHeart adeptly captures the intricate relationships between multiple heart components using a unified deformation field and introduces an implicit registration technique to manage the pose variability in medical data. Built on WHS3D dataset of 140 refined whole-heart structures, ImHeart delivers superior reconstruction accuracy and anatomical fidelity. Moreover, we demonstrate the ImHeart can significantly improve heart segmentation from multi-center MRI scans through a retraining pipeline, adeptly navigating the domain gaps inherent to such data.

**Keywords:** implicit fields · heart · shape modeling · multi-part structures · data-centric.

## 1 Introduction

Implicit functions have emerged as a potent tool for shape modeling across various fields, including CAD modeling [14,11,2] and medical imaging [27,29,19]. Compared to explicit shape modeling approaches, such as point clouds [17,28] and meshes [7,24], implicit / hybrid methods offer significant flexibility and scalability [26]. However, within the medical domain, the application of these techniques has largely been limited to modeling single-class shapes, without adequately addressing the interrelations among different structures.

In heart modeling, accurately representing the connections between structures such as the left ventricle (LV) and the left atrium (LA) is essential, with the LV myocardium ideally enveloping parts of the LV. This aspect has been significantly overlooked thus far. The straightforward use of single-class implicit

\* Correspondence to: Jiancheng Yang (jiancheng.yang@epfl.ch).

\*\* This author contributed during a research internship at EPFL.

functions to model these components in isolation often proves inadequate for capturing the complex interrelations among different heart parts, potentially leading to topological discrepancies. While some studies have addressed the issue of topological correctness in biomedical imaging [13,6] and implicit functions [10], their focus has predominantly been on analyzing (pairwise) connectivity. This approach becomes fragile for modeling the whole-heart structure, which involves up to eight (or more) types of structures with various topological relationships.

Inspired by template-based implicit methods [30,19,27], which model target shapes through deformed learnable templates, we recognize the utility of using learnable templates to implicitly capture relationships among multiple classes. Specifically, we characterize the multi-class relationship with a multi-class template and model it using a unified deformation field to maintain the relative relationships among multiple classes as much as possible. Furthermore, considering the variability in poses often encountered in medical data collection, we introduce a novel implicit registration technique that facilitates simple modifications for achieving registration within the implicit field. Further inspired by [30], we also incorporate an implicit residual to model the remaining shape components. While theoretically not precise, this practical strategy has proven effective. Compared to methods that come with theoretical guarantees, such as diffeomorphisms [19], we find relaxing these constraints can enhance the representational capability without increasing topological errors.

To conduct this study, we compiled a dataset of whole-heart structures, named *WHS3D*, consisting of 140 refined shapes. Based on *WHS3D*, we introduce *ImHeart*, a latent variable model designed to capture and deform the complex eight structures of the heart. This model builds upon the foundation of implicit deformable template methods [27]. Our findings demonstrate that *ImHeart* is capable of effectively representing intricate heart structures, achieving superior reconstruction metrics while maintaining anatomical precision.

As a practical application, we illustrate how *ImHeart*, in conjunction with an imperfect heart segmentation model, can enhance the quality of heart segmentations derived from multi-center cardiac MRI scans. We employ a potentially undertrained neural network to produce an initial segmentation; then, we fit *ImHeart* model to this initial output to refine the segmentation, and subsequently re-train the segmentation network using the refined output. This data-centric AI pipeline, yields anatomically accurate segmentations, serving as a crucial application. In a clinical setting, one rarely has enough multi-center training data to properly train a segmentation network to handle the quirks of the particular apparatus being used to acquire the images. This method makes it possible to bridge the domain gap these quirks introduce. Moreover, with high-quality annotations, it is compatible to integrate other techniques such as domain adaptation [5] and test-time training [20,22] to further improve the model performance.

## 2 Methods

### 2.1 Preliminaries

Implicit shape representations model shapes by assigning 3D coordinates to a shape indicator, typically choosing between occupancy or signed distance. In this work, we opt for the former for its simplicity. We employ linear interpolation, rendering our training occupancy field analogous to a truncated signed distance field. For a shape  $S$ , this mapping can be formally described as follows:

$$\mathcal{F}(\mathbf{h}, \mathbf{p}) = o : \mathbb{R}^c \times \mathbb{R}^3 \rightarrow \mathbb{R}, \quad (1)$$

where  $\mathbf{h}$  represents a  $c$ -dimensional *latent vector* that encodes the shape  $S$ ,  $\mathbf{p}(x, y, z) \in \mathbb{R}^3$  denotes a query point, and  $\mathcal{F}$ , implemented via a deep neural network, outputs the occupancy probability  $o \in [0, 1]$ . Given a set of training data, both  $\mathcal{F}$  and the vector  $\mathbf{h}$  corresponding to each shape can be learned through various methods [14,16]. In this study, we adopt an auto-decoding approach [14].

Beyond methods that directly generate shape indicators, there exist studies that conceptualize shapes as deformations from templates [3,30,19,27,25]. In these works, the function  $\mathcal{F}$  is a composite of two functions,  $\mathcal{T}$  and  $\mathcal{D}$ :

$$\mathcal{F}(\mathbf{h}, \mathbf{p}) = \mathcal{T}(\mathcal{D}(\mathbf{h}, \mathbf{p})), \quad (2)$$

where  $\mathcal{D} : \mathbb{R}^c \times \mathbb{R}^3 \rightarrow \mathbb{R}^3$  is a function that transforms a query point  $\mathbf{p}$  into new coordinates, and  $\mathcal{T} : \mathbb{R}^3 \rightarrow \mathbb{R}$  is a learned implicit function akin to  $\mathcal{F}$  in the original formulation but is specialized to learn a single shape. The functions  $\mathcal{D}$  and  $\mathcal{T}$  can be realized in various manners [19,27,30]. We use the approach where learning multiple templates is essential [27], we define  $\mathcal{T}$  as an *Implicit Template Network* and  $\mathcal{D}$  as an *Implicit Deformation Network*:

$$\mathcal{F}(\mathbf{h}, \mathbf{p}) = \mathcal{T}(\mathbf{t}, \mathbf{p} + \mathcal{D}(\mathbf{d}, \mathbf{p})), \quad (3)$$

with  $\mathbf{t}$  and  $\mathbf{d}$  being distinct vectors derived from  $\mathbf{h}$ . Notably,  $\mathcal{T}$  incorporates an additional parameter, enabling the selection of one among several potential templates using a straight-through estimator [1].

### 2.2 ImHeart: Implicit Field on Multi-Class Heart Structure

ImHeart is constructed based on the framework outlined in Eq. 3. However, we extend it to accommodate multi-class scenarios by applying a unified deformation to multi-class templates. Additionally, we introduced an implicit registration technique to address the variability in poses frequently observed in the collection of medical data. Furthermore, we incorporate an implicit residual component to capture the remaining shape elements that are not addressed by the deformation or template alone. The complete ImHeart model is depicted in Fig. 1. Given a latent vector  $\mathbf{h}$ , we define  $\mathbf{t} = \mathbf{t}(\mathbf{h})$ ,  $\mathbf{d} = \mathbf{d}(\mathbf{h})$ , and  $\mathbf{r} = \mathbf{r}(\mathbf{h})$  as distinct vectors derived from  $\mathbf{h}$ . These vectors are processed through their respective implicit networks— $\mathcal{T}$  for templates,  $\mathcal{D}$  for deformation, and  $\mathcal{R}$  for residuals. Following this setup, we provide a formal description of our ImHeart model.

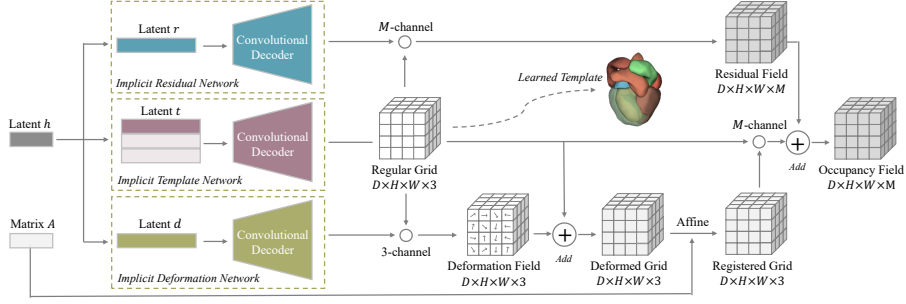


Fig. 1: **ImHeart**. Given a latent vector  $\mathbf{h}$ , we initiate the process by deriving three key vectors:  $\mathbf{t}$ ,  $\mathbf{d}$ , and  $\mathbf{r}$ . Each of these vectors is then processed through their respective implicit networks— $\mathcal{T}$  for templates,  $\mathcal{D}$  for deformation, and  $\mathcal{R}$  for residuals. To address the variability often observed in medical imaging due to different poses, an optimizable matrix  $\mathbf{A}$  is applied to the deformed grid.

*Unified Deformation on Multi-Class Templates.* We employ a unified deformation field  $\mathcal{D} : \mathbb{R}^c \times \mathbb{R}^3 \rightarrow \mathbb{R}^3$  applied to a multi-class template,  $\mathcal{T} : \mathbb{R}^c \times \mathbb{R}^3 \rightarrow \mathbb{R}^M$ , where  $M$  represents the number of classes of interest. In our case, focusing on 8 heart structures plus the background, we have  $M = 9$ . Through this approach, we model the relationships among multiple classes using templates and maintain these relationships via a unified deformation field.

*Implicit Registration.* To accommodate the variability in poses commonly seen in medical data collection, we introduce an implicit registration technique. This involves an affine transformation on the deformed coordinates, defined as:

$$\mathbf{p}' = \mathbf{R} \cdot (\mathbf{p} + \mathcal{D}(\mathbf{d}, \mathbf{p})) + \mathbf{b}, \quad (4)$$

where  $\mathbf{A} = (\mathbf{R}, \mathbf{b})$  with  $\mathbf{R} \in \mathbb{R}^{3 \times 3}$  being a freely optimized affine matrix and  $\mathbf{b} \in \mathbb{R}^3$  a shift vector. This means our method includes rotation, scaling and shearing. Per-sample  $\mathbf{A}$  can be optimized together with the corresponding  $\mathbf{h}$  using auto-decoding, with  $\mathbf{R}$  initially set to the identity matrix and  $\mathbf{b}$  to zero. There are also related idea in previous studies for neural rendering [15].

*Implicit Residual.* Drawing inspiration from prior work [30], we further integrate an implicit residual component through an Implicit Residual Network  $\mathcal{R} : \mathbb{R}^c \times \mathbb{R}^3 \rightarrow \mathbb{R}^M$ . This allows us to model the remaining shape components more accurately. The final model output is then given by:

$$\mathcal{F}(\mathbf{h}, \mathbf{p}) = \mathcal{T}(\mathbf{t}, \mathbf{p}') + \mathcal{R}(\mathbf{r}, \mathbf{p}), \quad (5)$$

ensuring a comprehensive representation of the complex heart structures by blending the template deformation with residual shape adjustments.

*Implementation Details.* To efficiently process 3D images, we implement coordinates through a  $D \times H \times W \times 3$  3D grid. On the architecture, we adopt a convolutional feature decoder inspired by the ConvONet [16], which leverages




Fig. 2: **Visualization of WHS3D Dataset.** We show both the source shapes using semi-automatic methods [31,4] and after our repairing.

the spatial hierarchies and local features inherent in 3D shapes. Regarding model training, considering our method is an extension of ImplicitAtlas [27], we maintain consistency with its training recipe.

### 2.3 Data-Centric Pipeline to Generate Annotations

We demonstrate the capability of ImHeart to improve the quality of heart segmentations obtained from multi-center cardiac MRI scans. The approach involves initially relying on a segmentation network trained on some data to obtain an initial segmentation on the target data. In real practice, due to the significant domain gaps and limited training data characteristic of multi-center data, the segmentation results from the initial network are often suboptimal. Subsequently, we refine these initial segmentations by optimizing the latent vector of ImHeart to reconstruct the initial segmentation. The reconstructed result is then used as a pseudo ground truth to re-train the segmentation network on the target data. Notably, this pipeline is data-centric, aiming to generate anatomically accurate annotations on the target dataset. Afterwards, it is still possible to integrate other techniques such as domain adaptation [5] and test-time training [20,22] to further improve the model performance.

## 3 Datasets

### 3.1 WHS3D: Multi-Class Shapes for Whole-Heart Structure

To train ImHeart, we compiled a high-quality heart shape dataset that encompasses 8 structures: myocardium of the left ventricle, left ventricle, right ventricle, left atrium, right atrium, pulmonary artery, ascending aorta, and descending aorta. The data originate from heart segmentations performed using semi-automatic methods [31,4] across multiple source datasets [21,18,12]. These annotations, while comprehensive, contained numerous artifacts.

To address the issue, we employed Neural Annotation Refinement [29], a technique designed to repair and refine annotation quality. Furthermore, we applied several empirical rules aimed at ensuring the topological accuracy of the labels. The rules include several morphological operations to ensure that the 8 classes each have only 1 connected component (CC). If not, we retain the




Fig. 3: Domain Gap between MMWHS-MRI [32] and In-House Data.

largest CC after a closing operation. Additionally, we ensure connectivity between classes: ventricle vs atrium, aorta vs left ventricle, and right ventricle vs pulmonary artery. Since these labels come from real data, disconnected cases are rare (usually due to annotation noise), and we slightly dilate to ensure proper connection. Fig. 2 visualizes some samples, including the source and repaired.<sup>3</sup>

### 3.2 Downstream Cardiac Segmentation Datasets

To validate the effectiveness of our proposed retraining pipeline based on ImHeart, we utilized two downstream cardiac segmentation datasets. MMWHS-MRI [32], is a widely used cardiac MRI dataset that includes all WHS3D structures except for the descending aorta. This dataset comprises a total of 20 annotated samples. We allocated 15 of these for the training set and reserved 5 for testing. Additionally, we employed an in-house dataset consisting of 15 cardiac MRI cases obtained from a clinical radiology department, which similarly encompasses the same seven labels as those present in the MMWHS-MRI.

Due to differences in imaging conditions, there exists a significant domain gap between the in-house data and the MMWHS-MRI dataset, as illustrated in Fig 3. This gap is particularly evident in the lower resolution and blurred edges of the in-house data, providing a substantial ground for testing the transferability and adaptability of the ImHeart model to real-world clinical data.

## 4 Experiments

### 4.1 Heart Shape Modeling

*Settings.* To evaluate the shape modeling performance of ImHeart on the WHS3D dataset, we adopted a setup similar to the conventional auto-decoding approach [14], dividing WHS3D into 120 known (training) and 20 unknown (test) samples. Our performance analysis focuses on the unknown shapes at a resolution of 128. Besides modeling on the eight structures of WHS3D as a whole, we also conducted a separate modeling for the myocardium (Myo) for analysis.

We utilized two sets of metrics for our analysis. Volumetric Dice Score (DSC): A standard metric ranging from 0 to 1, higher is better. We report the micro-average DSC across multiple classes. Number of Connected Components (CC):

<sup>3</sup> The WHS3D dataset is available at [cvlab.epfl.ch/data](http://cvlab.epfl.ch/data).

Table 1: **Heart Shape Reconstruction Results on WHS3D.** We compare with several baselines with or without templates. In our methods, Sep/Uni denotes separate/unified deformation, and IR denotes implicit registration.

Methods	Single-Class (Myo)		Multi-Class (Avg-8)	
	DSC $\uparrow$	CC $\downarrow$	DSC $\uparrow$	CC $\downarrow$
<i>w/o Template</i>				
DeepSDF [14]	0.866	1.15	0.863	1.74
NeAR [29]	0.913	<b>1.00</b>	0.927	1.59
<i>w/ Template</i>				
DIT [30]	0.899	<b>1.00</b>	0.915	1.21
NDF [19]	0.896	<b>1.00</b>	0.904	1.13
ImplicitAtlas [27]	0.904	<b>1.00</b>	0.919	1.23
ImHeart Sep		<b>1.00</b>	0.920	1.26
ImHeart Uni	<b>0.939</b>	<b>1.00</b>	0.941	1.04
ImHeart Uni+IR	0.938	<b>1.00</b>	<b>0.956</b>	<b>1.03</b>




Fig. 4: **Visualization of Generated Heart Samples by ImHeart.**

We design this metric to serve as a simple assessment of topological correctness. We calculate the average number of connected components across different classes, lower is better. The ideal value is 1, indicating that each class has only one connected component, which aligns with the expected anatomical reality. All results are averaged over three experiments.

As comparison, we analyzed a range of baseline methods, with or without templates. These baselines are trained separately on the multiple classes. Additionally, ImHeart was examined in various configurations, including separate (Sep) and unified (Uni) deformation fields, as well as implicit registration (IR).

*Results.* The results presented in Tab. 1 indicate that, generally, methods based on templates contribute to a reduction in CC, suggesting that templates, regardless of their specific design, benefit the topology of the modeled structures. It is noteworthy that NDF [19], despite its theoretical topological guarantees, does not perfectly address all issues in practice and has a relatively lower DSC.

Unified deformation significantly enhances the model performance in multi-class reconstruction, compared to the separate one ( $p=0.0127$ , t-test). Furthermore, IR can lead to additional improvements ( $p=0.0162$ , t-test), although IR does not have a substantial impact when modeling a single class.

Table 2: **Downstream Segmentation Results on MMWHS-MRI [32] and In-House Data.** DSC: Dice score. CC: number of connected components.

Methods	MMWHS-MRI (Avg-7)		In-House (Avg-7)	
	DSC $\uparrow$	CC $\downarrow$	DSC $\uparrow$	CC $\downarrow$
Initial Baseline	0.871	1.17	0.673	3.06
Naïve Retraining	0.864	1.09	0.635	2.70
ImHeart Retraining	<b>0.874</b>	<b>1.03</b>	<b>0.777</b>	<b>1.08</b>

In Fig. 4, we visualize several samples generated by ImHeart. These visualizations showcase the model capability to produce anatomically accurate samples.

#### 4.2 Downstream Application: Segmentation on Unseen Dataset

In Tab. 2, we report the performance on downstream data using the same evaluation metrics. Additionally, in Fig. 5, we visualize the 2 sample results.

We establish an initial baseline by training a model on the MMWHS-MRI dataset. Specifically, we use STU-Net [8] pretrained on TotalSegmentator [23], which adopts the training strategy of nnU-Net [9]. While this initial baseline demonstrates strong performance on the MMWHS-MRI test set (both in terms of DSC and CC), its direct application on the In-House data results in poor performance due to the large domain gap (as shown in Fig. 3).

Following this, we apply our proposed ImHeart Retraining pipeline to refine labels, on MMWHS-MRI and In-House datasets, respectively. It is observed that, even on the MMWHS-MRI dataset, there is still a slight enhancement, especially in the reduction of CC. On the In-House dataset, the improvements in both DSC and CC are significant compared to the initial baseline. The visualizations in Fig. 5 further illustrate that the cardiac structures generated using our method appear anatomically normal.

Notably, employing a Naïve Retraining strategy, which uses the predictions from the initial baseline without ImHeart refinement, does not improve performance; it may even degrade it. This highlights the value of integrating ImHeart into the retraining pipeline, as it provides anatomical knowledge that significantly enhance segmentation accuracy, particularly in overcoming challenges posed by domain gaps.

## 5 Conclusion

We introduce ImHeart, a novel approach to multi-class heart shape modeling by integrating learnable templates with a unified deformation field, addressing the complex interrelations among multiple heart structures. Developed on a dataset of 140 whole-heart structures, it achieves high accuracy and anatomical precision. Furthermore, ImHeart significantly enhances heart segmentation in multi-center MRI scans, overcoming domain gaps through a retraining pipeline.



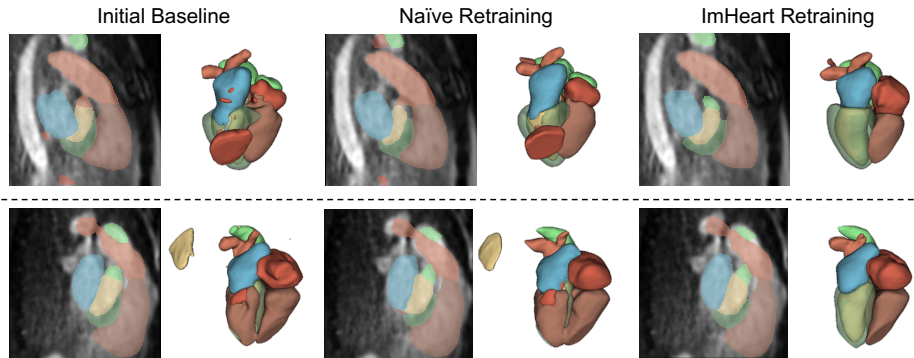


Fig. 5: Visualization of the Predicted Results of Initial Baseline, Naïve Retraining and ImHeart Retraining.

**Acknowledgment.** This work was supported by a Swiss National Science Foundation grant.

**Disclosure of Interests.** Jiancheng Yang holds equity in Dianei Technology but believes this does not constitute a competing interest. Other authors declare that they have no conflict of interest.

## References

1. Bengio, Y., Léonard, N., Courville, A.: Estimating or propagating gradients through stochastic neurons for conditional computation. arXiv Preprint (2013)
2. Chen, Z., Zhang, H.: Learning Implicit Fields for Generative Shape Modeling. In: Conference on Computer Vision and Pattern Recognition (2019)
3. Deng, Y., Yang, J., Tong, X.: Deformed implicit field: Modeling 3d shapes with learned dense correspondence. In: Conference on Computer Vision and Pattern Recognition. pp. 10286–10296 (2021)
4. Gao, S., Zhou, H., Gao, Y., Zhuang, X.: Bayeseg: Bayesian modeling for medical image segmentation with interpretable generalizability. *Medical Image Analysis* (2023)
5. Guan, H., Liu, M.: Domain adaptation for medical image analysis: a survey. *IEEE Transactions on Biomedical Engineering* **69**(3), 1173–1185 (2021)
6. Gupta, S., Hu, X., Kaan, J., Jin, M., Mpoy, M., Chung, K., Singh, G., Saltz, M., Kurc, T., Saltz, J., et al.: Learning topological interactions for multi-class medical image segmentation. In: European Conference on Computer Vision. pp. 701–718. Springer (2022)
7. Hanocka, R., Hertz, A., Fish, N., Giryas, R., Fleishman, S., Cohen-Or, D.: Meshcnn: a network with an edge. *ACM Transactions on Graphics* **38**(4), 1–12 (2019)
8. Huang, Z., Wang, H., Deng, Z., Ye, J., Su, Y., Sun, H., He, J., Gu, Y., Gu, L., Zhang, S., et al.: Stu-net: Scalable and transferable medical image segmentation models empowered by large-scale supervised pre-training. arXiv Preprint (2023)

9. Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature methods* **18**(2), 203–211 (2021)
10. Le, H., Talabot, N., Yang, J., Fua, P.: Enforcing topological interaction between implicit surfaces via uniform sampling. *arXiv Preprint* (2023)
11. Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy Networks: Learning 3D Reconstruction in Function Space. In: *Conference on Computer Vision and Pattern Recognition*. pp. 4460–4470 (2019)
12. Metz, C., Schaap, M., Weustink, A., Mollet, N., van Walsum, T., Niessen, W.: Coronary centerline extraction from ct coronary angiography images using a minimum cost path approach. *Medical physics* **36**(12), 5568–5579 (2009)
13. Mosinska, A., Marquez-Neila, P., Koziński, M., Fua, P.: Beyond the pixel-wise loss for topology-aware delineation. In: *Conference on Computer Vision and Pattern Recognition*. pp. 3136–3145 (2018)
14. Park, J.J., Florence, P., Straub, J., Newcombe, R.A., Lovegrove, S.: DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation. In: *Conference on Computer Vision and Pattern Recognition* (2019)
15. Park, K., Sinha, U., Barron, J.T., Bouaziz, S., Goldman, D.B., Seitz, S.M., Martin-Brualla, R.: Nerfies: Deformable neural radiance fields. In: *International Conference on Computer Vision*. pp. 5865–5874 (2021)
16. Peng, S., Niemeyer, M., Mescheder, L., Pollefeys, M., Geiger, A.: Convolutional occupancy networks. In: *European Conference on Computer Vision*. pp. 523–540. Springer (2020)
17. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: *Conference on Computer Vision and Pattern Recognition*. pp. 652–660 (2017)
18. Schaap, M., Metz, C.T., van Walsum, T., van der Giessen, A.G., Weustink, A.C., Mollet, N.R., Bauer, C., Bogunović, H., Castro, C., Deng, X., et al.: Standardized evaluation methodology and reference database for evaluating coronary artery centerline extraction algorithms. *Medical Image Analysis* **13**(5), 701–714 (2009)
19. Sun, S., Han, K., Kong, D., Tang, H., Yan, X., Xie, X.: Topology-preserving shape reconstruction and registration via neural diffeomorphic flow. In: *Conference on Computer Vision and Pattern Recognition*. pp. 20845–20855 (2022)
20. Sun, Y., Wang, X., Liu, Z., Miller, J., Efros, A., Hardt, M.: Test-time training with self-supervision for generalization under distribution shifts. In: *International Conference on Machine Learning*. pp. 9229–9248. PMLR (2020)
21. Tobon-Gomez, C., Geers, A.J., Peters, J., Weese, J., Pinto, K., Karim, R., Ammar, M., Daoudi, A., Margeta, J., Sandoval, Z., et al.: Benchmark for algorithms segmenting the left atrium from 3d ct and mri datasets. *IEEE Transactions on Medical Imaging* **34**(7), 1460–1473 (2015)
22. Tomar, D., Vray, G., Bozorgtabar, B., Thiran, J.P.: Tesla: Test-time self-learning with automatic adversarial augmentation. In: *Conference on Computer Vision and Pattern Recognition*. pp. 20341–20350 (2023)
23. Wasserthal, J., Breit, H.C., Meyer, M.T., Pradella, M., Hinck, D., Sauter, A.W., Heye, T., Boll, D.T., Cyriac, J., Yang, S., et al.: Totalsegmentator: Robust segmentation of 104 anatomic structures in ct images. *Radiology: Artificial Intelligence* **5**(5) (2023)
24. Wickramasinghe, U., Remelli, E., Knott, G., Fua, P.: Voxel2mesh: 3D Mesh Model Generation from Volumetric Data. In: *Conference on Medical Image Computing and Computer Assisted Intervention* (2020)

25. Wickramasinghe, U., Jensen, P., Shah, M., Yang, J., Fua, P.: Weakly supervised volumetric image segmentation with deformed templates. In: Conference on Medical Image Computing and Computer Assisted Intervention. pp. 422–432. Springer (2022)
26. Xie, K., Yang, J., Wei, D., Weng, Z., Fua, P.: Efficient anatomical labeling of pulmonary tree structures via implicit point-graph networks. arXiv Preprint (2023)
27. Yang, J., Wickramasinghe, U., Fua, P.: ImplicitAtlas: Learning Deformable Shape Templates in Medical Imaging. In: Conference on Computer Vision and Pattern Recognition (2022)
28. Yang, J., Gu, S., Wei, D., Pfister, H., Ni, B.: Ribseg dataset and strong point cloud baselines for rib segmentation from ct scans. In: Conference on Medical Image Computing and Computer Assisted Intervention. pp. 611–621. Springer (2021)
29. Yang, J., Shi, R., Wickramasinghe, U., Zhu, Q., Ni, B., Fua, P.: Neural annotation refinement: Development of a new 3d dataset for adrenal gland analysis. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 503–513. Springer (2022)
30. Zheng, Z., Yu, T., Dai, Q., Liu, Y.: Deep implicit templates for 3d shape representation. In: Conference on Computer Vision and Pattern Recognition. pp. 1429–1439 (2021)
31. Zhuang, X.: Multivariate mixture model for myocardial segmentation combining multi-source images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **41**(12), 2933–2946 (2018)
32. Zhuang, X., Li, L., Payer, C., Štern, D., Urschler, M., Heinrich, M.P., Oster, J., Wang, C., Smedby, Ö., Bian, C., et al.: Evaluation of algorithms for multi-modality whole heart segmentation: an open-access grand challenge. *Medical Image Analysis* **58**, 101537 (2019)