

Pairwise-Constrained Implicit Functions for 3D Human Heart Modeling

Hieu Le^{*1}, Jingyi Xu², Nicolas Talabot¹, Jiancheng Yang¹, and Pascal Fua¹

¹ EPFL, Lausanne, Switzerland

² Stony Brook University, New York, USA

{first-name,last-name}@epfl.ch, jingyixu@cs.stonybrook.edu

Abstract. Accurate 3D models of the human heart require not only correct outer surfaces but also realistic inner structures, such as the ventricles, atria, and myocardial layers. Approaches relying on implicit surfaces, such as signed distance functions (SDFs), are primarily designed for single watertight surfaces, making them ill-suited for multi-layered anatomical structures. They often produce gaps or overlaps in shared boundaries. Unsigned distance functions (UDFs) can model non-watertight geometries but are harder to optimize, while voxel-based methods are limited in resolution and struggle to produce smooth, anatomically realistic surfaces. We introduce a pairwise-constrained SDF approach that models the heart as a set of interdependent SDFs, each representing a distinct anatomical component. By enforcing proper contact between adjacent SDFs, we ensure that they form anatomically correct shared walls, preserving the internal structure of the heart and preventing overlaps, or unwanted gaps. Our method significantly improves inner structure accuracy over single-SDF, UDF-based, voxel-based, and segmentation-based reconstructions. We further demonstrate its generalizability by applying it to a vertebrae dataset, preventing unwanted contact between structures.

Keywords: Implicit Function · 3D Heart Reconstruction · Spine Reconstruction · Topology

1 Introduction

Accurate 3D modeling of the human heart is key to surgical planning, computational simulations, and disease diagnosis. However, the heart’s multi-chambered structure, with its ventricles, atria, and myocardial layers, poses a fundamental challenge: Accurate external reconstruction is not enough. Shared surfaces between individual components must be preserved, while preventing gaps, overlaps, and incorrect connections.

Despite recent advances, existing methods still fail to do this. Voxel-based approaches such as nn-UNet [7,8] are resolution-limited, making smooth surface

* corresponding author; email: minh.le@epfl.ch

modeling difficult. Classic implicit functions, such as SDFs [16], yield higher-resolution reconstructions but are best for single watertight surfaces. Newer SDF-based approaches [13,22] handle composite shapes but focus on overall shape reconstruction rather than delivering anatomical accuracy or meeting anatomical constraints. Unsigned distance functions (UDFs) [2,17] can model inner structures directly, but suffer from unstable optimization, which limits accuracy.

To address these challenges, we propose a novel pairwise-constrained implicit function approach that enforces proper anatomical contacts. We model the heart as a set of interdependent SDFs, each corresponding to a distinct anatomical component. Unlike previous methods that focus only on adjacency[4] or avoiding interpenetration[9], our approach ensures that adjacent components form contact regions with correct surface areas, accurately reflecting anatomical measurements. This matters because, when analyzing cardiac datasets, it has been observed that the ratio between the surface area of the contact region and the total surface area of paired components remain remarkably stable [1,15,3].

This yields a previously unused constraint that we leverage to improve reconstruction accuracy. To this end, we introduce a sampling-based verification strategy, where we randomly sample a large number of points in space to measure the actual contact ratio between components. By comparing this empirical ratio to the expected anatomical values from the training set, we iteratively adjust the implicit function representation, ensuring that the reconstructed heart maintains both topological and anatomical correctness.

We validate our method on cardiac MRI datasets, demonstrating that it significantly improves inner structure accuracy over single-SDF, UDF-based, and voxel-based approaches. To demonstrate its generalizability, we also apply it to a vertebrae dataset, where it successfully maintains proper spacing between adjacent vertebrae - a very different constraint from the heart. This highlights the promise of pairwise-constrained implicit functions for achieving high-fidelity, multi-component medical 3D modeling.

2 Related work

Modeling the internal structure of the human heart requires preserving anatomical relationships between adjacent chambers while ensuring accurate reconstruction of shared surfaces. Some segmentation-based methods enforce spatial consistency using topology-aware losses [5,6] or multi-category spatial constraints [4], but their voxelized nature limits resolution and smoothness. Mesh-based approaches like Voxel2Mesh [21] refine outputs into meshes but struggle with complex anatomical details due to their reliance on simple geometric templates.

Classic implicit functions, such as SDFs [16], yield high-resolution reconstructions but were only intended for single, watertight surfaces. Multi-SDF approaches such as MODIF [13], IPM[22], or SOMH [19] extend this to more complex structures but do not explicitly enforce inter-object constraints needed for accurate shared surface reconstruction. Our method addresses this limitation by enforcing anatomical contact regions with precise surface area constraints,

ensuring structural integrity. UDFs [2] provide an alternative for modeling non-watertight surfaces but suffer from unstable optimization [11,17,18], making them unreliable for medical applications. Our experiments show that, when applied to heart modeling, UDF-based methods often generate incomplete or missing surfaces, further limiting their usability.

3 Method

Our method ensures that adjacent structures maintain proper contacts of the correct extent. In heart reconstruction, it preserves the correct amount of the shared surface area between chambers [1]. These shared surface ratios are remarkably stable and can be reliably derived from the training set. For example, the wall between the left ventricle and its myocardium consistently accounts for 27% of their combined surface area and analysis of 120 training instances only shows deviations of less than 2% in the most extreme cases. Our approach enforces these ratios, preventing unnatural gaps or overlaps. Similarly, when applying our approach to the reconstruction of healthy human spines, where adjacent vertebrae must maintain a minimum gap of at least 1 mm [12], we introduce separation constraints.

Our framework handles both scenarios equally well, even though they involve very different constraints. For contact ratios, the primary difficulty lies in ensuring that objects interact without penetrating each other, which requires precise alignment and continuous adjustment of surface boundaries. For minimum distance constraints, detecting violations necessitates global checks to ensure that no parts of the objects come too close from each other. Addressing these two distinct sets of requirements using traditional modeling frameworks can be cumbersome and often involves separate, specialized approaches. In contrast, our approach handles both cases in a consistent, unified manner.

Our method relies on SDFs [16], which have emerged as a powerful model to learn continuous representations of 3D shapes. They allow detailed reconstructions of object instances as well as meaningful interpolations between them. Given an object, a signed distance function outputs the point’s distance to the closest object surface: $f(\mathbf{z}, \mathbf{x}) = s : \mathbf{x} \in \mathbb{R}^3, s \in \mathbb{R}$, where \mathbf{z} is a latent vector that parameterizes the surface. In the remainder of this section, we first introduce our approach to enforcing contact ratio constraint in the heart. We then show how it can be extended to enforcing the minimum distance constraint between vertebrae.

3.1 Enforcing Correct Shared Surface Areas

Given objects represented by their SDFs and a prior indicating a desired contact ratio (%) between them, our goal is to refine their 3D shapes in such a way that they do not intersect and touch each other within a small tolerance of the desired contact ratio. We do so by randomly sampling points and identifying ones lying in critical regions and adjusting the distances of all implicit shapes

to those points to achieve the desired contact ratio. The key questions we need to address are *how do we select these points*, *how many do we need*, and *how to use them to adjust the shapes*.

In our experiment, we reconstruct simultaneously five components of the human heart: The four ventricles and the myocardium of the left-ventricle, with the contact ratios between all pairs being given. However, for notational simplicity, in the remainder of this section we describe our approach in terms of only two components with SDFs, f_A and f_B , whose network weights are frozen. We refine the shapes by minimizing the losses defined below with respect to the corresponding latent vectors, \mathbf{a} and \mathbf{b} . This generalizes naturally to all five.

Mining Topologically Meaningful Points. To enforce shared surface constraints, we introduce loss functions for three key point sets obtained by random sampling: (1) $\mathcal{A}_{\text{contact}}$ – points that should be in contact with both objects, (2) $\mathcal{A}_{\text{intersecting}}$ – points that sit inside both objects, and (3) $\mathcal{A}_{\text{ncontact}}$ – points that should not be in contact areas. First, for two *explicit* surfaces A, B , the contact ratio between A and B is defined as

$$P_{A,B} = \frac{\text{Area}(S_{AB})}{\text{Area}(S_A) + \text{Area}(S_B)}, \quad (1)$$

where S_{AB} represents the partial surface of object A that is in close proximity, *i.e.*, contact distance, to object B , and S_A and S_B refer to the entire surfaces of objects A and B , respectively. For implicit objects, this can be approximated as

$$P'_{A,B} = \frac{\sum_{i=1}^N \mathbb{1}(|f_A(\mathbf{a}, \mathbf{x}_i)| < \epsilon) \cdot \mathbb{1}(|f_B(\mathbf{b}, \mathbf{x}_i)| < \epsilon)}{\sum_{i=1}^N \mathbb{1}(|f_A(\mathbf{a}, \mathbf{x}_i)| < \epsilon) + \sum_{i=1}^N \mathbb{1}(|f_B(\mathbf{b}, \mathbf{x}_i)| < \epsilon)}, \quad (2)$$

where $\{\mathbf{x}_i : i \in (1, N)\}$ is a set of N uniformly sampled random points, ϵ is a small threshold indicating the contact distance, and $\mathbb{1}(\cdot)$ is the indicator function that returns the value 1 if its statement is true and 0 otherwise.

Our goal is to adjust \mathbf{a} and \mathbf{b} so that the estimated contact surface $P'_{A,B}$ aligns with the expected value $P_{A,B}$. To achieve this, we modify the numerator of eq. (2), which represents the contact surface. While the denominator, which reflects the combined surface areas of the objects, could also be altered, it tends to remain stable during optimization as the overall object sizes are largely determined by the initial segmentation inputs. Thus, the expected number of points in the contact region should be

$$N_{\text{contact}} = P_{A,B} \times \left(\sum_{i=1}^N \mathbb{1}(|f_A(\mathbf{a}, \mathbf{x}_i)| < \epsilon) + \sum_{i=1}^N \mathbb{1}(|f_B(\mathbf{b}, \mathbf{x}_i)| < \epsilon) \right), \quad (3)$$

where $P_{A,B}$ is the expected proportion of shared surface area between the two components, which can be easily computed from the training shapes. To enforce this expected number of contact points, we define three sets of points for optimization. 1) $\mathcal{A}_{\text{contact}}$: The top N_{contact} points with closest summed distance to both objects. 2) $\mathcal{A}_{\text{ncontact}}$: All points that are in close proximity to both f_A and

f_B but are not included in $\mathcal{A}_{\text{contact}}$, and 3) $\mathcal{A}_{\text{intersecting}}$: all points that are within both objects: $\mathcal{A}_{\text{intersecting}} = \{\mathbf{x} \mid f_A(\mathbf{a}, \mathbf{x}) < 0 \wedge f_B(\mathbf{b}, \mathbf{x}) < 0\}$. These three sets, which are periodically resampled, allow us to define the following loss functions whose minimization ensures correct contacts, while preventing overlap.

Contact Ratio Loss. To ensure the correct shared surface, we define

$$\mathcal{L}_{\text{contact}} = \sum_{\mathbf{x} \in \mathcal{A}_{\text{contact}}} |f_A(\mathbf{a}, \mathbf{x}) + f_B(\mathbf{b}, \mathbf{x})| \text{ and } \mathcal{L}_{\text{ncontact}} = \sum_{\mathbf{x} \in \mathcal{A}_{\text{ncontact}}} -(f_A(\mathbf{a}, \mathbf{x}) + f_B(\mathbf{b}, \mathbf{x})). \quad (4)$$

Minimizing $\mathcal{L}_{\text{contact}}$ encourages contact points to lie equally between both surfaces without overlap, reaching zero when $f_A(\mathbf{a}, \mathbf{x}) = -f_B(\mathbf{b}, \mathbf{x})$. Minimizing $\mathcal{L}_{\text{ncontact}}$ prevents excessive contact.

Intersection and data Losses. To prevent intersections and ensure that objects remain faithful to their segmentation, we define

$$\mathcal{L}_{\text{intersecting}} = \sum_{\mathbf{x} \in \mathcal{A}_{\text{intersecting}}} |f_A(\mathbf{a}, \mathbf{x})| + |f_B(\mathbf{b}, \mathbf{x})|, \quad (5)$$

$$\mathcal{L}_{\text{data}} = \sum_{\mathbf{x} \in \mathbf{X}_A} |f_A(\mathbf{a}, \mathbf{x}) - s_{\mathbf{x}}| + \sum_{\mathbf{x} \in \mathbf{X}_B} |f_B(\mathbf{b}, \mathbf{x}) - s_{\mathbf{x}}|, \quad (6)$$

where \mathbf{X}_A and \mathbf{X}_B represent 500'000 spatial locations per object and are sampled more aggressively (90%) near the object surfaces.

Optimization. We reconstruct both objects by minimizing

$$\mathcal{L} = \lambda_1 \mathcal{L}_{\text{intersecting}} + \lambda_2 \mathcal{L}_{\text{contact}} + \lambda_3 \mathcal{L}_{\text{ncontact}} + \lambda_4 \mathcal{L}_{\text{data}}. \quad (7)$$

To adapt to shape changes, we resample 300K points every 10 iterations to update topologically meaningful points.

3.2 Enforcing Minimum Distance Constraints

Our method can also ensure a minimum separation between adjacent structures when direct contact is not anatomically valid, such as between vertebrae in the spine. To achieve this, we identify points where the sum of distances to both surfaces falls below a given threshold d , forming the violation set $\mathcal{A}_{\text{violation}} = \{\mathbf{x} \mid f_A(\mathbf{a}, \mathbf{x}) + f_B(\mathbf{b}, \mathbf{x}) < d\}$. We then define the composite loss

$$\mathcal{L}_{\text{min-distance}} = \lambda_1 \sum_{\mathbf{x} \in \mathcal{A}_{\text{violation}}} \max(0, d - (f_A(\mathbf{a}, \mathbf{x}) + f_B(\mathbf{b}, \mathbf{x}))) + \lambda_2 \mathcal{L}_{\text{data}}, \quad (8)$$

whose minimization fits the data while pushing the surfaces apart.




Fig. 1: The two rows show results from an in-distribution (ID) and an out-of-distribution (OOD) heart sample. We show transparent meshes to highlight contact areas. Red indicates penetration, while green denotes proper shared surfaces.

Table 1: **3D Heart Reconstruction.** We compare methods based on Chamfer distance ($\times 10^4$), area differences (Area), and normal consistency (NC) relative to the ground-truth meshes.

Method	In-Distribution			In-House Testing Data		
	CD	Area	NC	CD	Area	NC
nnU-Net[8]	1.5	0.89	68.2	5.0	1.7	87.6
SDF [16]	1.6	2.66	98.8	5.0	2.0	96.8
IPM [22]	10.7	4.14	98.7	10.4	4.4	95.2
SOMH [19]	1.9	0.81	99.2	7.2	1.9	96.8
UDF [2]	1.5	0.40	89.0	40.6	1.8	73.4
NVF [17]	1.6	0.49	99.3	4.8	1.7	81.3
Ours	1.5	0.13	99.4	4.1	1.2	95.3

4 Experiments

We validate our method by reconstructing human heart and lumbar spines. We jointly reconstruct their components while enforcing proper constraints between them by fitting deep implicit functions [16] to segmentation outputs from nnU-Net [7]. We compare against several baselines: (1) a standard SDF method [16] that fits each part individually, (2) a multi-SDF method based on template deformation [22], (3) a multi-SDF approach that employs a single perceptron to implicitly learn part compatibility [19], (4) a simple UDF-based method [2] for modeling non-watertight shapes, and (5) a recent, state-of-the-art UDF-like representation for modeling open surfaces [17].

3D Heart Reconstruction. We use our method to reconstruct five heart components: left ventricle (LV), myocardium of the left ventricle (M-LV), left atrium

(LA), right atrium (RA), and right ventricle (RV). We train an SDF auto-decoder [16] on a public dataset of 120 whole-heart models [23], setting aside 20 validation samples used for choosing hyper-parameters. Contact ratios between components are precomputed from the training data and used as constraints during reconstruction. The nnU-Net segmentation model is trained on 15 cardiac images from the same dataset, with only 20 images publicly available. We evaluate our method on two test sets: (1) a public set of five cardiac images from the same distribution as the training data, and (2) an in-house dataset of 10 cardiac images from a local hospital, serving as an out-of-distribution test set due to significant differences in image quality. The ground-truth meshes are produced by applying Marching-Cube[14] and Laplacian smoothing on the annotated segmentations, similar to [10]. We assess reconstruction accuracy using Chamfer distance (CD), surface area distance (Area), and normal consistency (NC).

Table 2: **Contact Ratio Statistics.** Our method can output shapes with similar contact ratios observed from the training set.

Pairs	LV-MLV	MLV-RV	LV-RA	RV-LA
Training Data	27.0% \pm 0.06	8.6% \pm 0.03	5.5% \pm 0.01	7.6% \pm 0.03
Our output (Val)	27.1% \pm 0.01	8.8% \pm 0.01	5.4% \pm 0.03	7.3% \pm 0.01
Our output (Test)	25.5% \pm 0.02	10.0% \pm 0.08	5.0% \pm 0.03	5.6% \pm 0.08




Fig. 2: **Vertebrae reconstruction.** The second row shows zoomed-in crops of the first row. Yellow indicates close but allowable proximity, while green marks undesired contact.

As shown in Tab. 1, our method outperforms existing approaches across both datasets. For in-distribution data, it achieves the lowest CD (1.5, tied with nnU-Net and UDF), the lowest area difference (0.13), and the highest NC (99.4%). Area difference is crucial as it indicates the absence of double or missing surfaces, and our method performs best in this regard. On in-house data, we maintain the best CD and area difference while achieving strong NC. Fig. 1 further illustrates these results. For the OOD sample in the bottom row, nnU-Net struggles with severe misalignments and fragmented structures, as seen in the second row. In-

dividually fitting SDFs reduces Chamfer distance by imposing priors on noisy data but introduces major topological errors. Our method further lowers CD while preserving structural integrity, outputting shapes with very similar contact ratios to ones in the training set, as can be seen in Tab. 2. While SOMH[19] achieves high NC, this is primarily due to overly smoothed outputs, as it fails to find an embedding that properly fits the noisy input segmentation. UDF-based methods [2,17] do not produce any surface penetration by design, but their reconstructed meshes exhibit many missing holes due to unstable optimization.

Table 3: **Spine reconstruction.** Comparison of three methods by constraint violations (area & point count) and Chamfer distance (CD) across vertebrae pairs, and overall averages.

Method	L1 - L2		L2 - L3		L3 - L4		L4 - L5		All
	px ²	N	px ²	N	px ²	N	px ²	N	CD ($\times 10^3$)
nn-Unet	261.4	123.2	527.2	237.3	644.6	294.3	761.1	358.9	0.4
SDF	315.8	154.2	309.8	166.1	416.4	227.7	503.3	279.4	0.4
Ours	0.0	0.8	0.5	2.3	0.5	2.9	1.2	3.9	0.3

3D Lumbar Spine Segmentation. We also use our method to reconstruct vertebrae in a dataset containing 460 CT images of the five lower vertebrae of the human spine (L1-L5) [20]. Of these, 80% are used to train both the nnU-Net and the latent implicit models, while 10% samples are reserved for testing and 10% are for validation. In the entire dataset, each pair of adjacent vertebrae exhibits a minimum gap of 1 pixel, which is the constraint we enforce during reconstruction. In Tab. 3, we report topological errors measured by the number of contact vertices (N) and the contact surface areas (px²), along with reconstruction accuracy. Our approach yields reconstructions with almost no constraint violations. Note that enforcing the minimum gap constraint is particularly challenging because there are only small volumes surrounding each vertebrae joints that are topologically relevant, as can be seen in the highlighted areas in Fig. 2.

Table 4: **Ablation study.** Performance of our method when reconstructing a left-ventricle and its myocardium when removing specific loss components.

Method	w/o L_{inter}	w/o $L_{contact}$	w/o $L_{non-contact}$	w/o L_{data}	All	GT
Penetration	5%	0%	0%	0%	0%	0%
Contact	17.2%	13%	29%	26%	27%	27%
CD ($\times 10^4$)	3.8	3.7	3.5	33.2	3.1	0

Ablation. We gauge the effect of each loss function of Section 3.1 by omitting each one in turn when reconstructing a left-ventricle / myocardium pair in the out-of-distribution test set. As can be seen by comparing the “w/o” columns of Tab. 4 against the one that shows the results when using “all” losses, they all contribute to the accuracy of the final result.

5 Conclusion

We have introduced a novel method to enforcing topological consistency constraints between multiple 3D objects modeled in terms of implicit surfaces. In the heart reconstruction scenario, our method maintains a precise contact ratio while preventing interpenetrations. In the lumbar spine reconstruction case, we ensure that adjacent vertebrae does not touch each other. This is achieved via randomized sampling of well chosen object locations. Future work will explore more complex scenarios, such as enforcing the constraints only at specific locations and not at others where they do not apply due to various pathologies. This will be turned into a powerful diagnostic tools to detect the extent of these pathologies.

Acknowledgments. This work was supported in part by the Swiss National Science Foundation.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Buckberg, G., Nanda, N.C., Nguyen, C., Kocica, M.J.: What is the heart? anatomy, function, pathophysiology, and misconceptions. *Journal of Cardiovascular Development and Disease* **5** (2018), <https://api.semanticscholar.org/CorpusID:46934645>
2. Chibane, J., Mir, A., Pons-Moll, G.: Neural Unsigned Distance Fields for Implicit Function Learning. In: *Advances in Neural Information Processing Systems* (2020)
3. Choy, J.S., Kassab, G.S.: Scaling of myocardial mass to flow and morphometry of coronary arteries. *Journal of applied physiology* **104** **5** (2008), <https://api.semanticscholar.org/CorpusID:1085066>
4. 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)
5. Hu, X., Li, F., Samaras, D., Chen, C.: Topology-Preserving Deep Image Segmentation. In: *Advances in Neural Information Processing Systems*. pp. 5658–5669 (2019)
6. Hu, X., Wang, Y., Fuxin, L., Samaras, D., Chen, C.: Topology-Aware Segmentation Using Discrete Morse Theory. In: *International Conference on Learning Representations* (2021)
7. Isensee, F., Petersen, J., Zimmerer, A.K., Jaeger, P., Kohl, S., Wasserthal, J.: NnuNet: Self-Adapting Framework for U-Net-Based Medical Image Segmentation. In: *arXiv Preprint* (2018)
8. Isensee, F., Wald, T., Ulrich, C., Baumgartner, M., Roy, S., Maier-Hein, K.H., Jaeger, P.F.: nnu-net revisited: A call for rigorous validation in 3d medical image segmentation. *ArXiv* (2024)
9. Karunratanakul, K., Yang, J., Zhang, Y., Black, M., Muandet, K., Tang, S.: Grasping Field: Learning Implicit Representations for Human Grasps. In: *International Conference on 3D Vision* (2020)
10. Kong, F., Wilson, N., Shadden, S.C.: A Deep-Learning Approach for Direct Whole-Heart Mesh Reconstruction. *Medical Image Analysis* **74** (2021)
11. Le, H., Stella, F., Guillard, B., Fua, P.: Gradient distance function (2024), <https://arxiv.org/abs/2410.22422>
12. Little, J.S., Khalsa, P.S.: Material properties of the human lumbar facet joint capsule. *Journal of biomechanical engineering* **127** **1**, 15–24 (2005), <https://api.semanticscholar.org/CorpusID:25597804>
13. Liu, Y., Planche, B., Zheng, M., Gao, Z., Sibut-Bourde, P., Yang, F., Chen, T., Wu, Z.: Implicit modeling of non-rigid objects with cross-category signals. In: *AAAI Conference on Artificial Intelligence* (2024)
14. Lorensen, W., Cline, H.: Marching Cubes: A High Resolution 3D Surface Construction Algorithm. In: *ACM SIGGRAPH*. pp. 163–169 (1987)
15. Medrano-Gracia, P., Cowan, B.R., Ambale-Venkatesh, B., Bluemke, D.A., Eng, J., Finn, J.P., Fonseca, C.G., Lima, J.A.C., Suinesiaputra, A., Young, A.A.: Left ventricular shape variation in asymptomatic populations: the multi-ethnic study of atherosclerosis. *Journal of Cardiovascular Magnetic Resonance* **16**, 56 – 56 (2014), <https://api.semanticscholar.org/CorpusID:12885556>
16. 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)

17. Rella, E.M., Chhatkuli, A., Konukoglu, E., Gool, L.V.: Neural vector fields for implicit surface representation and inference. *International Journal of Computer Vision* (2024)
18. Stella, F., Talabot, N., Le, H., Fua, P.: Neural Surface Localization for Unsigned Distance Fields. In: *European Conference on Computer Vision* (2024)
19. Verhülsdonk, J., Grandits, T., Costabal, F.S., Pinetz, T., Krause, R., Auricchio, A., Haase, G., Pezzuto, S., Effland, A.: Shape of my heart: Cardiac models through learned signed distance functions. In: *Proceedings of The 7nd International Conference on Medical Imaging with Deep Learning. Proceedings of Machine Learning Research*, vol. 250. PMLR (03–05 Jul 2024)
20. Wasserthal, J., Breit, H.C., Meyer, M.T., Pradella, M., Hinck, D., Sauter, A.W., Heye, T., Boll, D., Cyriac, J., Yang, S., Bach, M., Segeroth, M.: Totalsegmentator: Robust segmentation of 104 anatomic structures in ct images. *Radiology. Artificial intelligence* **5** **5** (2022)
21. 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)
22. Zhang, C., Elgharib, M., Fox, G., Gu, M., Theobalt, C., Wang, W.: An implicit parametric morphable dental model. *ACM Trans. Graph.* (2022)
23. 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)