



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.

Algebraic Sphere Surface Fitting for Accurate and Efficient Mesh Reconstruction from Cine CMR Images

Jin He¹, Weizhou Liu¹, Shifeng Zhao¹, Yun Tian^{1(✉)}, and Shuo Wang²

¹ School of Artificial Intelligence, Beijing Normal University, Beijing, China
tianyun@bnu.edu.cn

² Digital Medical Research Center, School of Basic Medical Sciences, Fudan University, Shanghai, China

Abstract. Accurate 3D modeling of the ventricles through cine cardiovascular magnetic resonance (CMR) imaging benefits precise clinical assessment of cardiac morphology and motion. However, the existing short-axis stacks exhibit low spatial resolution in the inter-slice orientation compared to the intra-slice direction, resulting in a sparse representation of the realistic heart. The anisotropic short-axis images pose challenges in directly reconstructing meshes from them. In this work, we propose a surface fitting approach based on the algebraic sphere, which serves as a previous step for various mesh-based applications, to reconstruct a natural ventricular shape from the segmented wireframe-type point cloud. Considering the sparse and layered nature of the point clouds, we first estimate the normals of the point cloud based on dynamic programming and neighborhood selection, followed by fitting a point set surface using a non-compact kernel adapted by layers. Finally, an implicit scalar field representing the signed distance between the query point and the projection point is obtained, and the manifold mesh is extracted by meshing zero iso-surface. Experimental results on two publicly available datasets demonstrate that the proposed framework can accurately and effectively reconstruct ventricular mesh from a single image with better cross-domain generalizability.

Keywords: Wireframe-type point clouds · Mesh reconstruction · Algebraic Sphere · Point set surfaces · Cine CMR.

1 Introduction

Cine imaging is considered the primary approach among the imaging techniques available in CMR. The imaging sequence employed in this study enables precise visualization of vascular structures and facilitates the detection of ventricular motion abnormalities. An accurate and holistic comprehension of cardiac structure and function is made possible by the 3D visualization of the heart. However, rapid scanning of 3D+time short-axis stacks generally achieves high temporal

resolution while exhibiting anisotropic spatial resolution, leading to reduced resolution in the out-of-plane orientation (Fig. 1a), making it difficult to accurately and efficiently reconstruct the shape of the heart.

In the field of medical imaging, there has been recent interest in integrating meshes to evaluate ventricular structures or functions. A common approach [1, 3, 5, 10, 26] involves creating template meshes and adjusting the weights of statistical shape models (SSMs) to match each case in the dataset. However, generating SSMs requires learning weights, and SSM-based methods may not fully align with the contours of 2D images. Another approach [2, 6, 12, 17, 25, 26] involves utilizing neural networks to learn vertex deformations of template meshes. Kong et al.’s method [12, 13] is more suitable for isotropic CT data. Meng et al. [17] employed high-resolution segmentation and linear interpolation to obtain ground truth meshes, but their approach is limited to specific populations and not applicable to diseased data. Additionally, high-resolution segmentation [18, 20, 22] can be used to describe the 3D shape of the heart, but it can only generate images at fixed resolutions. Deep learning-based methods often require a number of high-resolution images or rely on meshes as ground truth. However, ground truth meshes are typically obtained through image interpolation (Fig. 1d), super-resolution segmentation with post-processing methods such as Marching cubes [16] (Fig. 1c), or derived from SSMs.




Fig. 1. (a) Cine images offer only a sparse representation of the actual 3D heart. (b) Proficient graphics method [9] for wireframe-type point clouds still falls short in surface reconstruction. (c) Directly reconstructed meshes using Marching Cubes lack guaranteed topology due to anisotropic data. (d) Interpolating information in each layer results in poor mesh quality, with post-processing often leads to mesh shrinkage and distortion. (e) Our method achieves accurate and efficient reconstruction of smooth manifold meshes.

Our main focus is on the preliminary stage of mesh-based tasks, such as the aforementioned segmentation, generation, and motion estimation, particularly on the direct reconstruction of surface mesh from an individual CMR image. We perceive the contour derived from the segmented image as a distinctive wireframe-type point cloud, thereby formulating the problem as the reconstruction of surface mesh from wireframe. Decoupling segmentation and reconstruction offers enhanced interpretability and seamless integration of cutting-edge

segmentation methods. Recently, extensive research has been conducted in the field of computer graphics on surface reconstruction from point clouds. However, when dealing with the wireframe-type data which exhibits sparsity and hierarchical characteristics, both deep learning-based [14, 23, 24] and geometry-based [8, 9, 15] approaches yield unsatisfactory outcomes (Fig. 1b). To address this shortcoming, we propose a novel method based on a key insight: the cardiac surface is a smooth manifold. Leveraging this understanding, we develop a fitting method based on algebraic sphere to rapidly reconstruct a smooth manifold surface from the point cloud extracted from a single cine CMR image.

Our contribution has three folds. First, we tackle the challenge of reconstructing mesh from an individual CMR image without requiring datasets or training efficiently and accurately. Second, due to the special wireframe-type data, we propose an algebraic sphere surface fitting framework and introduce a layered adaptive non-compact kernel and a normal estimation method. Third, our method outperforms existing method and demonstrates excellent generalization performance.

2 Methodology

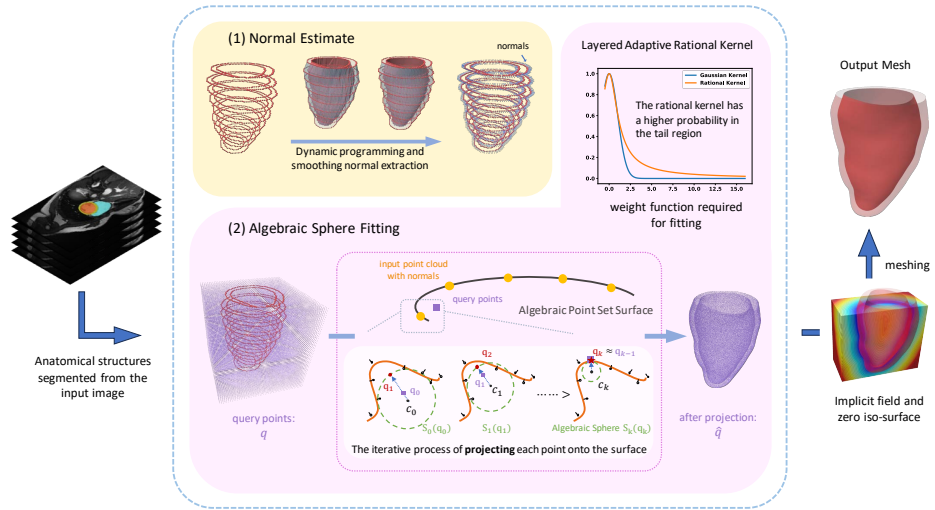


Fig. 2. An overview of the proposed method. (1) To reliably drive the algebraic sphere surfaces fitting, dynamic programming is employed to achieve rough meshing and Each point in the point cloud calculates a principal normal based on its neighborhood. (2) With algebraic spheres as local fitting primitives for projection using a rational kernel as the weight function and the Algebraic Point Set Surfaces (APSS) [7] serving as the underlying surface to recover large missing data, the proposed method extracts a precise mesh from the zero iso-surface in the implicit field.

Extracting point cloud $\mathcal{P} = \{(\mathbf{p}_i, l_i) | p_i \in \mathbb{R}^3, l_i \in \{1, 2, \dots, L\}\}_{i=1}^N$ from a 3D cine image, our method outputs mesh M representing the reconstructed heart, where N is the number of points and L is the number of layers. We denote the point cloud normals as $\{\mathbf{n}_i\}_{i=1}^N$, which need to be estimated.

Algebraic Sphere Fitting. The main advantage of the sphere fitting is its significantly improved stability in situations where planar moving least squares fails. In contrast to geometric sphere fitting, algebraic spheres allows us to elegantly handle planar areas or regions around inflection points as limits in which the algebraic sphere naturally degenerates to a plane.

As is shown in Fig. 2, given a query point $\mathbf{q} \in \mathbb{R}^3$, an algebraic sphere $S : \mathbb{R}^3 \rightarrow \mathbb{R}$ and projects \mathbf{q} is continuously optimized onto the zero iso-surface of S . Here the algebraic sphere S is a scalar field in \mathbb{R}^3 and its algebraic expression is $S(\mathbf{x}) := u_0 + \mathbf{u}_{123}\mathbf{x} + u_4\mathbf{x}^\top\mathbf{x}$ to \mathcal{P} , where $u_0 \in \mathbb{R}$, $\mathbf{u}_{123} \in \mathbb{R}^3$ and $u_4 \in \mathbb{R}$ are respectively the constant, linear and quadratic coefficients of the sphere. For simplicity, we denote $\mathbf{u} = (u_0, \mathbf{u}_{123}, u_4)$, and for a point \mathbf{q} , $\mathbf{u}(\mathbf{q})$ is the solution to the following optimization problem:

$$\mathbf{u}(\mathbf{q}) = \arg \min_{\mathbf{u}(\mathbf{q})} \sum_{i=1}^N \mathcal{H}(\mathbf{p}_i, \mathbf{q}) \cdot (S(\mathbf{p}_i)^2 + \beta \|\nabla S(\mathbf{p}_i) - \mathbf{n}_i\|^2) \quad (1)$$

where $\mathcal{H}(\cdot, \cdot)$ is kernel function and β is a constant default 10^9 .

Projection Process. The projection operation relies solely on local point clouds, which enhances the scalability and enables parallel acceleration during surface reconstruction. To retrieve the actual surface, one need to iteratively project the query points onto the zero iso-surface of current algebraic sphere. This process can be described as follows:

$$(\mathbf{q}_0 = \mathbf{q}, S_0 = S) \xrightarrow{\mathcal{R}_1} (\mathbf{q}_1, S_1) \xrightarrow{\mathcal{R}_2} (\mathbf{q}_2, S_2) \xrightarrow{\mathcal{R}_3} \dots \xrightarrow{\mathcal{R}_k} (\mathbf{q}_k, S_k) \quad (2)$$

where \mathcal{R}_i projects \mathbf{q}_{i-1} onto the zero iso-surface of S_{i-1} , denoted as $\mathcal{R}_i : \mathbf{q}_i = \mathbf{proj}(\mathbf{q}_{i-1}, S_{i-1})$, S_i is the algebraic sphere associated with \mathbf{q}_i , \mathbf{q}_k and S_k represent the projected point and algebraic sphere at the convergence of the projection process. In practical scenarios, convergence to the surface can be achieved with k as low as 10. Assuming the point $\hat{\mathbf{q}}$ is the projection of \mathbf{q} onto the zero iso-surface of S , $\mathbf{proj}(\mathbf{q}, S)$ is represented as follows:

$$\hat{\mathbf{q}} = \mathbf{proj}(\mathbf{q}, S) = \begin{cases} \mathbf{q} - \frac{S(\mathbf{q})}{\|\nabla S(\mathbf{q})\|} \nabla S(\mathbf{q}), & \text{if } u_4 = 0 \\ \mathbf{q} - (r - \|\mathbf{q} - \mathbf{c}\|) \frac{\nabla S(\mathbf{q})}{\|\nabla S(\mathbf{q})\|}, & \text{otherwise} \end{cases} \quad (3)$$

where $\mathbf{c} = -\frac{1}{2u_4} \cdot \mathbf{u}_{123}$ and $r = \sqrt{\mathbf{c}^\top \mathbf{c} - u_0/u_4}$ represent the center and radius of S respectively.

Finally, one can extract a surface by contouring the implicit function $f(\mathbf{q}) = (\mathbf{q} - \mathbf{q}_k) \cdot \mathbf{n}_{\mathbf{q}_k}$, where $\mathbf{n}_{\mathbf{q}_k} = \frac{\nabla S(\mathbf{q}_k)}{\|\nabla S(\mathbf{q}_k)\|}$. In practice, we use uniformly sampled grid (Fig. 2.(2)) as query points for the purpose of meshing using the marching cubes. The code is available at <https://github.com/hejin9/algebraic-sphere-surface-fitting>.

Layered Adaptive Rational Kernel. The kernel function $\mathcal{H}(\cdot, \cdot)$ plays a crucial role in the reconstruction results as it reflects the weight of each point in \mathcal{P} with respect to the query point \mathbf{q} . In previous works, commonly used kernels include the Gaussian kernel $\mathcal{H}(\mathbf{x}, \mathbf{y}) = e^{-\frac{\|\mathbf{x}-\mathbf{y}\|^2}{\lambda}}$ (which has compact support and is suitable for dense data but fails for non-uniform sparse point cloud data) and the rational kernel $\mathcal{H}(\mathbf{x}, \mathbf{y}) = (\|\mathbf{x} - \mathbf{y}\|^2 - 1)^2$ (non-compact support, used for sparse data). Although these methods perform well on typical data, they often result in a planar reconstruction for the same layer of point cloud, making them unable to handle the sparse and hierarchical point clouds extracted from CMR (as shown in Fig. 1b, where even with the addition of long-axis point clouds, a more accurate surface cannot be obtained). To address this, we have customized a layered adaptive rational kernel specifically tailored for the data: $\mathcal{H}(\mathbf{p}_i, \mathbf{q}) = (\|\mathbf{p}_i - \mathbf{q}\|^2 + \epsilon)^{-\frac{(l_i - l_c - 1)^2 + 2}{2}}$, where l_c represents the layer label of the point in \mathcal{P} that is closest to the point \mathbf{q} and ϵ is default 0.01. The kernel function we designed combines the characteristics of distance and data anisotropy, taking into account both the euclidean distance from the query point to \mathcal{P} and the layered neighborhood features.

Normal Estimate. The normal estimate is another key aspect, as the addition of normal makes algebraic-sphere-based fittings more robust. Due to the anisotropic nature of point clouds, traditional normal estimation methods often fail to achieve globally consistent and accurate normals, as referenced in the supplementary materials. Our coarse-to-fine approach begins by using dynamic programming to generate a roughly shaped triangular mesh and estimate the approximate normals of the point cloud.

Meshing adjacent layers for the endocardium and epicardium of the left ventricle is an independent and similar problem. Therefore, it is sufficient to design a dynamic programming specifically for the endocardium of the left ventricle. For two adjacent layers, with the number of points on the endocardium being N_U and N_D respectively, let's assume that each layer is sorted in counterclockwise order, denoted as $\{\mathbf{p}_i^U\}_{i=1}^{N_U}$ and $\{\mathbf{p}_i^D\}_{i=1}^{N_D}$ respectively and points \mathbf{p}_1^U and \mathbf{p}_1^D represent the closest pair of points between the two layers. Due to the numerous possibilities for triangulating the point clouds of the two layers, considering the Delaunay criterion, we aim to minimize the total edge length of the triangulated mesh. We design the states for dynamic programming as follows: $\mathcal{F}(i, j)$ represents the minimum total edge length of the triangulated mesh when connecting \mathbf{p}_i^U and \mathbf{p}_j^D with an edge. The state transition is as follows:

$$\mathcal{F}(i, j) = \min\{\mathcal{F}(i-1, j) + \|\mathbf{p}_{i-1}^U - \mathbf{p}_j^D\|, \mathcal{F}(i, j-1) + \|\mathbf{p}_i^U - \mathbf{p}_{j-1}^D\|\} \quad (4)$$

, the initial state is $\mathcal{F}(1, 1) = \|\mathbf{p}_1^U - \mathbf{p}_1^D\|$, and the terminal state is $\mathcal{F}(N_U, N_D)$. By backtracking from the terminal state, we can obtain the triangulated mesh of the adjacent layers (note that this coarse mesh lacks proper topology). We apply Laplace smoothing to the mesh to better incorporate neighborhood features and prior knowledge: the heart is a smooth manifold. The normals of the point cloud are then determined by the normals of the nearest vertices on the mesh.

3 Experiments

In this section, we use Automated Cardiac Diagnosis Challenge (ACDC) dataset [4] and Multi-Modality Whole Heart Segmentation (MMWHS) dataset [27] to benchmark our results. There are several experiments to evaluate the performance of the proposed method. First, the accuracy in reconstructing 3D meshes is validated using synthetic data based on the MMWHS dataset. Secondly, the ACDC dataset is used to evaluate the performance on real data. Lastly, the generalizability of the proposed method across various scenarios are explored.

3.1 Performance Evaluation of 3D Mesh Reconstruction

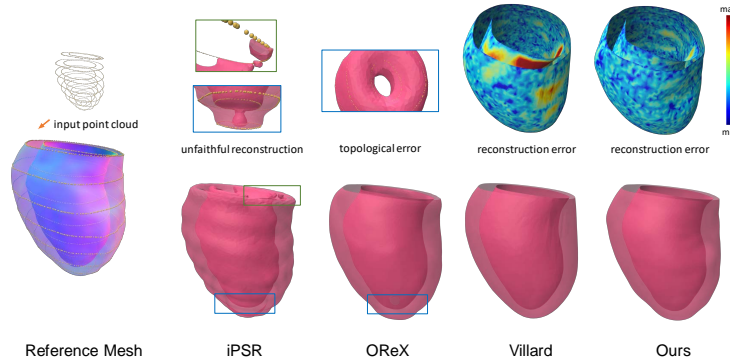


Fig. 3. Visualized meshes reconstructed from the synthetic data based on MMWHS dataset. The bottom row represents the reconstruction results.

To evaluate the accuracy of 3D mesh reconstruction, the MMWHS dataset is utilized for its high-resolution and isotropic properties. We extracted point clouds from the dataset, estimated normal consistency, and performed Poisson Surface Reconstruction (PSR) [11] to obtain a manifold reference mesh (Fig. 3). To simulate real cine images, layered wireframe-type point clouds are extracted as input. For comparison, we implemented three distinct methods in our study: iPSR [8], a geometry-based method; OReX [19], a learning-based method; and a myocardial mesh reconstruction method that we have named Villard [21]. iPSR

is an improved PSR, eliminating the need for point normals and employing an iterative approach. OReX, on the other hand, focuses on 3D shape reconstruction solely from slices and incorporates a neural field as the interpolation prior. Villard, through iterative mesh deformation, Laplacian smoothing, and subdivision, achieves the reconstruction of ventricular shapes.

Table 1. Comparison of other 3D mesh reconstruction methods based on simulated degradation to assess the similarity between reconstructed meshes and reference meshes using Chamfer Distance (CD), Hausdorff Distance (HD), Earth Mover’s Distance (EMD), and Average Symmetric Surface Distance (ASSD) metrics (mean \pm std). The Dice Similarity Coefficient(DSC) were employed to measure the similarity between the reconstructed meshes and input 2D contours per layer. Top performance of each column are in bold.

	iPSR	OReX	Villard	Ours
HD \downarrow	18.268 \pm 9.422	11.352 \pm 4.626	3.710 \pm 1.791	3.710 \pm 1.458
CD \downarrow	4.841 \pm 1.643	2.169 \pm 0.763	1.859 \pm 0.687	1.701 \pm 0.636
EMD \downarrow	6.786 \pm 2.201	3.261 \pm 1.448	1.194 \pm 0.399	1.189 \pm 0.533
ASSD \downarrow	2.446 \pm 0.817	1.089 \pm 0.384	0.930 \pm 0.344	0.850 \pm 0.318
DSC \uparrow	0.673 \pm 0.461	0.992 \pm 0.004	0.974 \pm 0.011	0.988 \pm 0.005

Table 1 presents the quantitative evaluation results for 3D myocardial surface reconstruction. Our method achieved the best performance in most metrics. Direct application of iPSR, which doesn’t excel in processing partial missing point clouds, did not yield satisfactory results in our specific application scenario. OReX demonstrated the best performance in 2D similarity metrics (Table 1 and Table 2) as it faithfully preserved the original data. However, as depicted in Fig. 3, the predicted 3D shape tends to exhibit topological errors such as holes. This is attributed to OReX relying on signed distance fields for interpolation without extrapolation capabilities. Villard showed similar quantitative and qualitative results to our method. Nevertheless, as Villard is based on mesh deformation, significant reconstruction errors were observed, particularly in regions where curvature varies greatly, as depicted in Fig. 3. Additionally, as revealed in the third experiment, our method demonstrated suitability for a wider range of data types compared to Villard.

3.2 Performance Evaluation on Real Data

The ACDC dataset was utilized to assess the performance on real-world data. Since the dataset lacks mesh data and the first experiment already evaluated the 3D metrics quantitatively and qualitatively, we solely assessed the alignment of the reconstructed mesh with the ground truth segmentation data on 2D slices. Table 2 presents the quantitative results of myocardial surface reconstruction on the real dataset, where our method showcases the best overall performance. As mentioned earlier, OReX remains faithful to the original data but may exhibit

topological errors in the 3D shape. Additionally, as evident from Table 2, generating a data with OReX remains time-consuming, whereas our method is much faster in comparison.

Table 2. Real-world performance on ACDC training dataset. We present three metrics (mean \pm std): DSC, HD, runtime (seconds). **Bold** indicates the best performance and the underline indicates the second best results.

	iPSR	OReX	Villard	Ours
DSC \uparrow	0.774 \pm 0.413	0.990 \pm 0.005	0.982 \pm 0.006	<u>0.984 \pm 0.007</u>
HD \downarrow	11.545 \pm 20.919	0.447 \pm 0.150	1.184 \pm 0.163	<u>1.156 \pm 0.356</u>
Time(s) \downarrow	29.452	4428.292	<u>10.379</u>	8.485

3.3 Generalization Analysis

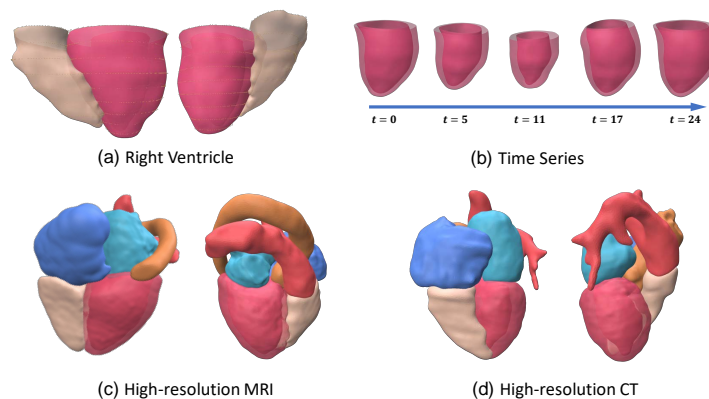


Fig. 4. Visualization of the versatility of our approach.

Fig. 4 demonstrates the versatility of our method, showcasing its applicability to a broader range of scenarios. In fact, in addition to reconstructing the left ventricle, we can also reconstruct the right ventricle, as shown in Fig. 4(a). For cardiac motion tracking, we can visualize the entire sequence by directly reconstructing mesh for each time frame. Furthermore, our method performs well with images of varying resolutions, allowing us to reconstruct both MRI and CT images. Additionally, it can generate complex shapes with ease. Furthermore, the input, being a point cloud and lacking explicit topological connectivity, allows our approach to adapt to various scenarios. As a result, our algorithm can be directly used without fine-tuning any parameters for other chambers, cine, multi-view inputs, and different disease morphologies. The details corresponding to Fig. 3 and Fig. 4 are included in the supplementary material.

4 Conclusion

To tackle the challenge of reconstructing 3D shapes from cine images, we proposed an innovative algebraic sphere surface fitting method that facilitates efficient and accurate mesh reconstruction. It is suitable for varying scenarios and achieves high performance on both simulated and real-world datasets. The reconstruction of dense and accurate geometric shapes of the heart for each case enables precise detection of myocardial locations exhibiting motion abnormalities, satisfying the visualization and interpretability requirements in the medical domain, and fostering further investigations into biomechanical analysis.

Acknowledgements. This work was supported by the National Natural Science Foundation of China(Grant No. 62172047) and the Beijing Natural Science Foundation(Grant No.4232027).

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

References

1. Attar, R., Pereañez, M., Bowles, C., Piechnik, S.K., Neubauer, S., Petersen, S.E., Frangi, A.F.: 3d cardiac shape prediction with deep neural networks: Simultaneous use of images and patient metadata. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2019. pp. 586–594 (2019)
2. Beetz, M., Banerjee, A., Grau, V.: Reconstructing 3d cardiac anatomies from misaligned multi-view magnetic resonance images with mesh deformation u-nets. In: Proceedings of the First International Workshop on Geometric Deep Learning in Medical Image Analysis. Proceedings of Machine Learning Research, vol. 194, pp. 3–14 (2022)
3. Beetz, M., Banerjee, A., Ossenbeng-Engels, J., Grau, V.: Multi-class point cloud completion networks for 3d cardiac anatomy reconstruction from cine magnetic resonance images. *Medical Image Analysis* **90**, 102975 (2023)
4. Bernard, O., Lalande, A., Zotti, C., Cervenansky, F., Yang, X., Heng, P.A., Cetin, I., Lekadir, K., Camara, O., Gonzalez Ballester, M.A., Sanroma, G., Napel, S., Petersen, S., Tziritas, G., Grinias, E., Khened, M., Kollerathu, V.A., Krishnamurthi, G., Rohé, M.M., Pennec, X., Sermesant, M., Isensee, F., Jäger, P., Maier-Hein, K.H., Full, P.M., Wolf, I., Engelhardt, S., Baumgartner, C.F., Koch, L.M., Wolterink, J.M., Išgum, I., Jang, Y., Hong, Y., Patravali, J., Jain, S., Humbert, O., Jodoin, P.M.: 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. Chen, X., Ravikumar, N., Xia, Y., Attar, R., Diaz-Pinto, A., Piechnik, S.K., Neubauer, S., Petersen, S.E., Frangi, A.F.: Shape registration with learned deformations for 3d shape reconstruction from sparse and incomplete point clouds. *Medical Image Analysis* **74**, 102228 (2021)
6. Deng, Y., Xu, H., Rodrigo, S., Williams, S.E., Williams, M.C., Niederer, S.A., Pushparajah, K., Young, A.: Modusgraph: Automated 3d and 4d mesh model reconstruction from cine cmr with improved accuracy and efficiency. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2023. pp. 173–183 (2023)
7. Guennebaud, G., Gross, M.: Algebraic point set surfaces. In: ACM SIGGRAPH 2007 Papers. SIGGRAPH '07 (2007)
8. Hou, F., Wang, C., Wang, W., Qin, H., Qian, C., He, Y.: Iterative poisson surface reconstruction (ipsr) for unoriented points. *ACM Transactions on Graphics* **41**(4) (2022)
9. Huang, Z., Carr, N., Ju, T.: Variational implicit point set surfaces. *ACM Transactions on Graphics* **38**, 1–13 (2019)
10. Joyce, T., Buoso, S., Stoeck, C.T., Kozerke, S.: Rapid inference of personalised left-ventricular meshes by deformation-based differentiable mesh voxelization. *Medical Image Analysis* **79**, 102445 (2022)
11. Kazhdan, M., Bolitho, M., Hoppe, H.: Poisson surface reconstruction. In: Proceedings of the Fourth Eurographics Symposium on Geometry Processing. p. 61–70. SGP '06 (2006)
12. Kong, F., Shadden, S.C.: Learning whole heart mesh generation from patient images for computational simulations. *IEEE Transactions on Medical Imaging* **42**(2), 533–545 (2023)
13. Kong, F., Wilson, N., Shadden, S.: A deep-learning approach for direct whole-heart mesh reconstruction. *Medical Image Analysis* **74**, 102222 (2021)

14. Li, Q., Feng, H., Shi, K., Gao, Y., Fang, Y., Liu, Y.S., Han, Z.: NeuralGF: Unsupervised point normal estimation by learning neural gradient function. In: Thirty-seventh Conference on Neural Information Processing Systems (NeurIPS) (2023)
15. Lin, S., Xiao, D., Shi, Z., Wang, B.: Surface reconstruction from point clouds without normals by parametrizing the gauss formula. *ACM Transactions on Graphics* **42**(2) (2022)
16. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3d surface construction algorithm. In: Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques. p. 163–169. SIGGRAPH '87 (1987)
17. Meng, Q., Bai, W., O'Regan, D.P., Rueckert, D.: Deepmesh: Mesh-based cardiac motion tracking using deep learning. *IEEE Transactions on Medical Imaging* pp. 1–1 (2023)
18. Sander, J., de Vos, B.D., Bruns, S., Planken, N., Viergever, M.A., Leiner, T., Išgum, I.: Reconstruction and completion of high-resolution 3d cardiac shapes using anisotropic cmri segmentations and continuous implicit neural representations. *Computers in Biology and Medicine* **164**, 107266 (2023)
19. Sawdayee, H., Vaxman, A., Bermano, A.H.: Ores: Object reconstruction from planar cross-sections using neural fields. In: 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 20854–20862 (2023)
20. Stolt-Ansó, N., McGinnis, J., Pan, J., Hammernik, K., Rueckert, D.: Nisf: Neural implicit segmentation functions. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 734–744. Springer (2023)
21. Villard, B., Grau, V., Zacur, E.: Surface mesh reconstruction from cardiac mri contours. *Journal of Imaging* **4**(1) (2018)
22. Wang, S., Qin, C., Savioli, N., Chen, C., O'Regan, D.P., Cook, S., Guo, Y., Rueckert, D., Bai, W.: Joint motion correction and super resolution for cardiac segmentation via latent optimisation. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2021. pp. 14–24 (2021)
23. Wang, Z., Wang, P., Wang, P., Dong, Q., Gao, J., Chen, S., Xin, S., Tu, C., Wang, W.: Neural-impls: Self-supervised implicit moving least-squares network for surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics* pp. 1–16 (2023)
24. Wang, Z., Zhang, Y., Xu, R., Zhang, F., Wang, P.S., Chen, S., Xin, S., Wang, W., Tu, C.: Neural-singular-hessian: Implicit neural representation of unoriented point clouds by enforcing singular hessian. *ACM Transactions on Graphics* **42**(6) (dec 2023)
25. Xia, Y., Chen, X., Ravikumar, N., Kelly, C., Attar, R., Aung, N., Neubauer, S., Petersen, S.E., Frangi, A.F.: Automatic 3d+t four-chamber cmr quantification of the uk biobank: integrating imaging and non-imaging data priors at scale. *Medical Image Analysis* **80**, 102498 (2022)
26. Yuan, X., Liu, C., Wang, Y.: 4d myocardium reconstruction with decoupled motion and shape model. In: 2023 IEEE/CVF International Conference on Computer Vision (ICCV). pp. 21195–21205 (2023)
27. Zhuang, X., Li, L., Payer, C., Štern, D., Urschler, M., Heinrich, M.P., Oster, J., Wang, C., Örjan Smedby, Bian, C., Yang, X., Heng, P.A., Mortazi, A., Bagci, U., Yang, G., Sun, C., Galisot, G., Ramel, J.Y., Brouard, T., Tong, Q., Si, W., Liao, X., Zeng, G., Shi, Z., Zheng, G., Wang, C., MacGillivray, T., Newby, D., Rhode, K., Ourselin, S., Mohiaddin, R., Keegan, J., Firmin, D., Yang, G.: Evaluation of algorithms for multi-modality whole heart segmentation: An open-access grand challenge. *Medical Image Analysis* **58**, 101537 (2019)