

HeartVolMesh: Cardiac Volumetric Mesh Reconstruction via Covariance-Guided Graph Deformation

Fengming Lin^{1,2,3}, Arezoo Zakeri^{1,2,4}, Haoran Dou^{1,2,3}, Zherui Zhou^{1,2},
Shaokun Lan^{1,2,3}, Jinming Duan^{1,2,4}, and Alejandro Frangi^{1,2,3,4,5*}

¹ Centre for Computational Imaging and Modelling in Medicine (CIMIM), University of Manchester, Manchester, UK

² Christabel Pankhurst Institute, University of Manchester, Manchester, UK

³ Department of Computer Science, University of Manchester, Manchester, UK

⁴ Division of Informatics, Imaging & Data Sciences, University of Manchester, Manchester, UK

⁵ NIHR Manchester Biomedical Research Centre, Manchester Academic Health Sciences Centre, University of Manchester, Manchester, UK

alejandro.frangi@manchester.ac.uk

<https://fmlinks.github.io/projects/heartvolmesh/>

<https://github.com/ccmim/HeartVolMesh>

Abstract. Accurate patient-specific tetrahedral cardiac meshes are essential for in-silico trials, yet common segmentation-then-modelling pipelines can blur thin-wall anatomy and offer limited cross-case correspondence. We propose HeartVolMesh, which lifts each template vertex to an anisotropic Gaussian kernel and uses a 3D CNN-GNN to predict per-vertex displacements and Cholesky-parameterized covariances from volumetric images. Training is guided by a covariance-aware negative log-likelihood loss with lightweight mesh regularization. For volumetric meshing, we warp a fixed tetrahedral template to the reconstructed surface via staged alignment, non-rigid registration, and deformation propagation, preserving connectivity and correspondence by construction, with resolution controlled by template density. Experiments show consistent gains over deformation-based baselines in surface mesh accuracy and volumetric mesh fidelity.

Keywords: Cardiac digital twins · Image-to-mesh learning · Graph neural networks · Volumetric meshing

1 Introduction

In-silico trials [1–3] increasingly evaluate cardiovascular devices and treatments using cardiac digital twins [5], reducing reliance on costly clinical studies [4, 3]. For FEM/CFD solvers, such twins require not only accurate anatomy but also directly usable 3D volumetric, typically tetrahedral, meshes [7, 6]. Automatically

* Corresponding author.




Fig. 1. Schematic illustration of different representations.

reconstructing topology-consistent, detail-faithful, high-quality cardiac volumetric meshes [28] from 3D medical images such as CTA is therefore central to scalable digital twins and in-silico trials.

Existing pipelines are still bottlenecked by volumetric meshing. Segmentation-then-meshing workflows use voxel-wise segmentation [18], marching-cubes extraction, and heavy post-processing [8, 9, 7]; however, simulation-ready tetrahedra still rely on heuristic surface-to-volume tetrahedralization, e.g., TetGen [10, 19, 20]. Such tools require case-specific tuning and yield case-dependent sampling/connectivity, limiting cross-case correspondence, population analysis [27], and standardized learning/simulation workflows.

End-to-end image-to-mesh methods [11, 12] reduce staging, and deformation-based models such as Voxel2Mesh, MeshDeformNet, and HeartDeformNet [11–13] preserve template topology for consistent surfaces. However, two main gaps remain: they use deterministic Euclidean supervision, which inadequately models boundary uncertainty near multi-structure junctions, and they remain surface-only, requiring heuristic tetrahedralization that breaks cross-case volumetric correspondence.

Motivated by these gaps, we seek an intermediate surface representation that captures fine details while explicitly accounting for vertex distribution, and a topology-preserving route to high-quality volumetric meshes with correspondence. We propose HeartVolMesh, a framework for cardiac digital twins with two complementary components:

1. We propose a covariance-guided graph deformation model that represents each template vertex as an anisotropic Gaussian kernel and couples vertices through the mesh graph, enabling covariance-aware and topology-consistent surface reconstruction.
2. We propose a volume-to-surface registration strategy with deformation-field propagation that warps a high-quality tetrahedral template to the patient-specific reconstructed surface, yielding tetrahedral meshes with correspondence and resolution controlled by template density.
3. We validate in multi-structure / whole-heart settings and demonstrate improved surface reconstruction accuracy, as well as new volumetric reconstruction capability that produces tetrahedral meshes with consistent cross-case vertex correspondence, together with better volumetric mesh quality across 24 cardiac structures.




Fig. 2. Overview of HeartVolMesh. From a CTA volume (Input 1), covariance-guided graph deformation reconstructs a topology-consistent target surface using anisotropic Gaussian vertices. A tetrahedral template mesh (Input 2) is then warped to this surface via volume-to-surface registration and deformation-field propagation, producing a patient-specific volumetric mesh with preserved connectivity and cross-case correspondence (Output).

2 Methodology

Our goal is to reconstruct topology-consistent and detail-faithful cardiac volumetric meshes from 3D medical images. The proposed HeartVolMesh framework consists of two complementary modules. First, we introduce covariance-guided graph deformation to reconstruct a topology-consistent surface while lifting each vertex to an anisotropic Gaussian kernel with a learned covariance. Second, we generate a tetrahedral mesh via a template-driven volume-to-surface registration strategy with deformation-field propagation, which warps a high-quality tetrahedral template to match the reconstructed surface and preserves cross-case vertex correspondence. During training, we supervise the surface module with a covariance-guided negative log-likelihood (NLL) matching loss and add lightweight mesh regularization to prevent degenerate deformations.

2.1 Covariance-guided Graph Deformation

Theory and Representation. Conventional deformation-based image-to-mesh methods typically predict only vertex displacements and enforce supervision with Euclidean distances (e.g., Chamfer) [26]. In medical images, however, the acceptable geometric tolerance is not uniform across space or directions [21] due

to partial-volume effects, imaging noise, and thin-wall anatomy. Our covariance-guided graph deformation addresses this by promoting each mesh vertex from a deterministic point to an anisotropic Gaussian kernel. Concretely, besides predicting the deformed vertex position (Gaussian mean), we predict a full 3×3 covariance matrix Σ per vertex, which extends a classical mesh representation $[N, 3], [E]$ to $[N, 3+6], [E]$, where E denotes mesh edges (graph connectivity). The covariance provides direction-dependent local geometric support, while the mesh graph couples neighboring vertices to preserve global anatomical coherence and topology consistency during deformation.

To ensure that Σ is always symmetric positive definite (SPD) and numerically stable, we parameterize it via a Cholesky factor $\Sigma = LL^\top + \epsilon I$. Here L is a lower-triangular matrix whose diagonal is constrained to be positive (we use a softplus transform), I is the identity matrix, and ϵ is a small constant that improves stability. This parameterization guarantees that Σ^{-1} and $\log \det(\Sigma)$ are well-defined for the covariance-guided loss described in Sec. 2.3.

Network Architecture. Fig. 2 shows the covariance-guided graph deformation network. A 3D CNN encoder [22] extracts multi-scale volumetric features, which are interpolated at current vertex locations to obtain image-conditioned vertex features. A GNN then propagates these features over the mesh graph, using neighborhood connectivity to enforce spatial coherence and topology-consistent deformation. We use coarse-to-fine refinement: meshes are progressively unpooled and updated across stages to capture both global structure and local detail. The final head predicts vertex displacements and covariance parameters for Cholesky-based Σ construction. In whole-heart settings, the encoder can be shared, with structure-specific mesh branches or output heads.

2.2 Template-driven Volumetric Mesh Generation

The bottom half of Fig. 2 summarizes our volume-to-surface registration pipeline. Let $\mathcal{S}$ be the template boundary surface extracted from a tetrahedral template, and $\mathcal{T}$ be the target surface reconstructed by covariance-guided graph deformation. We first perform a global similarity alignment [23] to initialize registration: $x' = sRx + t$, where $x \in \mathcal{S}$, R is rotation, t is translation, and s is a scale factor. We then apply a smooth non-rigid surface registration [24] to obtain per-vertex displacements on $\mathcal{S}$.

To propagate surface motion to the full volumetric template, we convert these surface displacements into a continuous 3D deformation field [25] $D(\cdot)$ defined on a voxel grid (via rasterization and interpolation/upsampling), and optionally smooth it (Gaussian blur) for regularity and stability. For multi-structure settings, we compute deformation fields per structure and combine them into a single dense field on the shared grid (write-to-position followed by interpolation and smoothing), consistent with Fig. 2. Finally, we warp all vertices of the tetrahedral template: $v_T = v_S + D(v_S)$, where v_S is a template volumetric vertex and v_T is the deformed vertex. This yields a patient-specific tetrahedral mesh

whose boundary matches the reconstructed surface while preserving template connectivity, correspondence, and element quality to a large extent.

2.3 Loss Function

We train the covariance-guided graph deformation module with a covariance-guided geometric matching objective plus lightweight mesh regularization. The volumetric warping module is executed at inference time as a registration-based optimization guided by the reconstructed surface.

Covariance-guided NLL Matching. For a matched pair of points $(x, y) \in \mathbb{R}^3$ with predicted covariance Σ , we define:

$$\ell(x, y; \Sigma) = \frac{1}{2}(x - y)^\top \Sigma^{-1}(x - y) + \beta \log \det(\Sigma). \quad (1)$$

The first term is a Mahalanobis distance under the learned anisotropic metric, and the second term prevents trivial covariance inflation. We use bidirectional nearest-neighbor matching (prediction-to-target and target-to-prediction) and average $\ell(\cdot)$ over matched pairs to form the surface loss L_{surf} .

Regularization and Overall Objective. To discourage degenerate deformations, we add standard mesh regularizers with small weights:

$$L = L_{\text{surf}} + \lambda_{\text{edge}} L_{\text{edge}} + \lambda_{\text{lap}} L_{\text{lap}} + \lambda_{\text{norm}} L_{\text{norm}}. \quad (2)$$

Here L_{edge} is edge-length regularization [16], L_{lap} is Laplacian smoothing [15, 17], and L_{norm} is normal consistency [14].

3 Experiments and Results

We evaluate whole-heart, multi-structure cardiac mesh reconstruction and compare with deformation-based image-to-mesh baselines. We report accuracy on both *pre-warp surfaces* (network output) and *post-warp volumetric meshes* (tetrahedral template warped to the predicted surface). Unless stated otherwise, results are mean $\pm$ std over the validation set, computed per structure.

3.1 Experimental Setup

We used a multi-centre in-house dataset of 900 patients and 4,000 temporal instances, split at the patient level into 800/100 patients for training/validation (3,600/400 instances). Reference pseudo-GT meshes were obtained from TotalSegmentator outputs with manual refinement, surface extraction, and the same non-learned template-warping procedure as in Sec. 2.2. All models were implemented in PyTorch and trained on one NVIDIA A100 80GB GPU. The 3D CNN encoder and mesh GNN were jointly optimized with Adam for 100 epochs,




Fig. 3. Qualitative volumetric reconstruction. Boundary surfaces extracted from the final tetrahedral meshes are visualized for LA/LV/RA/RV/Myo. Ours3p6 produces sharper boundaries and fewer outliers, especially near inter-structure junctions.

batch size 1, using 128^3 crops. The deformation field resolution was 300^3 . We used $\beta = 0.5$ and $\lambda_{edge} = \lambda_{lap} = \lambda_{norm} = 0.1$. Template-driven volumetric warping was applied after surface prediction and was not jointly trained. We evaluated surface reconstruction on LA, LV, RA, RV, and LVMyo, and volumetric reconstruction on LAMyo, LVMyo, RAMyo, and RVMyo using boundary surfaces extracted from the warped tetrahedral meshes. Metrics included CD, HD95, NC, inverted elements, minimum scaled Jacobian, and minimum dihedral angle. We compared Voxel2Mesh, MeshDeformNet, HeartDeformNet, Ours3p1, and Ours3p6, applying the same warping pipeline to all methods for fair volumetric evaluation.

3.2 Comparison with State-of-the-Art Methods

Volumetric meshes. Table 1 (left) evaluates the boundary surfaces extracted from the final warped tetrahedral meshes, measuring how accurately the resulting volumetric meshes align to the target anatomy. Our Ours3p6 achieves the best CD/HD95 and the highest NC across all four volumetric targets, demonstrating consistent improvements on both the thin-wall atrial myocardium (LAMyo) and the highly curved ventricular regions (LVMyo/RVMyo). While Ours3p1 already outperforms prior work, Ours3p6 further benefits from anisotropic covariance modelling and covariance-guided supervision, leading to more accurate and better-aligned volumetric boundaries. Beyond boundary accuracy, we also report

Table 1. Left: volumetric meshes (boundary surfaces). Right: surface meshes. Metrics: CD/HD95 in mm (lower is better) and NC in % (higher is better). Mean $\pm$ std. Bold indicates statistically significant improvement (t-test $p < 0.05$).

Category	Volumetric mesh (boundary surfaces)				Surface mesh				
	CD $\downarrow$								
Method	LVMyo	RVMyo	LAMyo	RAMyo	LA	LV	Myo	RA	RV
Voxel2Mesh	5.9 $\pm$ 0.5	4.3 $\pm$ 0.7	4.6 $\pm$ 1.5	3.7 $\pm$ 0.9	2.7 $\pm$ 0.2	3.5 $\pm$ 0.3	4.9 $\pm$ 0.4	3.5 $\pm$ 0.3	4.1 $\pm$ 0.5
MeshDeformNet	5.8 $\pm$ 0.5	4.4 $\pm$ 0.5	4.2 $\pm$ 1.4	4.3 $\pm$ 0.7	2.8 $\pm$ 0.2	3.7 $\pm$ 0.4	5.4 $\pm$ 0.5	3.9 $\pm$ 0.4	4.2 $\pm$ 0.8
HeartDeformNet	5.8 $\pm$ 0.5	4.5 $\pm$ 0.6	4.3 $\pm$ 1.6	4.3 $\pm$ 0.8	2.8 $\pm$ 0.2	3.7 $\pm$ 0.4	5.2 $\pm$ 0.5	3.7 $\pm$ 0.5	4.1 $\pm$ 0.6
Ours 3p1	4.5 $\pm$ 1.0	3.9 $\pm$ 0.9	4.0 $\pm$ 1.2	3.6 $\pm$ 0.9	2.6 $\pm$ 0.2	3.0 $\pm$ 0.4	4.1 $\pm$ 0.4	3.4 $\pm$ 0.3	3.8 $\pm$ 0.4
Ours 3p6	3.6$\pm$0.4	3.1$\pm$0.4	3.0 $\pm$ 0.3	3.0 $\pm$ 0.4	2.5$\pm$0.2	3.0$\pm$0.4	4.0$\pm$0.4	2.9$\pm$0.3	3.3$\pm$0.4
Ablation-A	5.8 $\pm$ 0.5	4.5 $\pm$ 0.6	4.3 $\pm$ 1.6	4.3 $\pm$ 0.8	2.8 $\pm$ 0.2	3.9 $\pm$ 0.5	5.3 $\pm$ 0.6	3.8 $\pm$ 0.6	4.5 $\pm$ 0.9
Ablation-B	3.9 $\pm$ 0.5	3.4 $\pm$ 0.5	2.9 $\pm$ 0.4	3.0 $\pm$ 0.3	/	/	/	/	/
Ablation-C	5.1 $\pm$ 0.9	4.5 $\pm$ 0.6	6.6 $\pm$ 1.4	6.4 $\pm$ 1.4	3.1 $\pm$ 0.2	3.4 $\pm$ 0.4	4.3 $\pm$ 0.4	3.3 $\pm$ 0.4	4.0 $\pm$ 0.4
	HD95 $\downarrow$								
Method	LVMyo	RVMyo	LAMyo	RAMyo	LA	LV	Myo	RA	RV
Voxel2Mesh	5.8 $\pm$ 0.7	4.8 $\pm$ 1.2	5.2 $\pm$ 1.7	4.2 $\pm$ 1.3	2.4 $\pm$ 0.2	3.3 $\pm$ 0.5	4.8 $\pm$ 0.4	3.8 $\pm$ 0.7	3.9 $\pm$ 0.7
MeshDeformNet	5.4 $\pm$ 0.6	4.9 $\pm$ 1.0	4.7 $\pm$ 1.5	4.7 $\pm$ 1.0	2.5 $\pm$ 0.2	3.6 $\pm$ 0.5	5.4 $\pm$ 0.6	4.5 $\pm$ 0.8	4.6 $\pm$ 2.0
HeartDeformNet	5.4 $\pm$ 0.7	4.9 $\pm$ 1.1	4.8 $\pm$ 1.7	4.7 $\pm$ 1.1	2.5 $\pm$ 0.2	3.8 $\pm$ 0.8	5.4 $\pm$ 0.7	4.2 $\pm$ 1.2	4.5 $\pm$ 1.6
Ours 3p1	4.2 $\pm$ 1.2	3.8 $\pm$ 1.3	4.4 $\pm$ 1.5	3.9 $\pm$ 1.3	2.3 $\pm$ 0.1	2.8 $\pm$ 0.4	3.8 $\pm$ 0.4	3.7 $\pm$ 1.0	3.7 $\pm$ 0.8
Ours 3p6	3.2$\pm$0.4	2.8$\pm$0.3	3.0 $\pm$ 0.5	2.9$\pm$0.6	2.2$\pm$0.2	2.8$\pm$0.5	3.6$\pm$0.3	2.8$\pm$0.5	3.1$\pm$0.5
Ablation-A	5.4 $\pm$ 0.7	5.0 $\pm$ 1.1	4.8 $\pm$ 1.7	4.8 $\pm$ 1.1	2.6 $\pm$ 0.3	3.8 $\pm$ 0.7	5.8 $\pm$ 0.7	4.4 $\pm$ 1.2	5.1 $\pm$ 2.3
Ablation-B	3.7 $\pm$ 0.4	3.1 $\pm$ 0.5	2.8 $\pm$ 0.4	3.0 $\pm$ 0.5	/	/	/	/	/
Ablation-C	5.0 $\pm$ 1.1	5.2 $\pm$ 1.0	7.7 $\pm$ 1.9	7.5 $\pm$ 1.4	2.8 $\pm$ 0.1	3.2 $\pm$ 0.5	4.0 $\pm$ 0.4	3.5 $\pm$ 0.8	4.0 $\pm$ 0.8
	Normal Consistency (%) $\uparrow$								
Method	LVMyo	RVMyo	LAMyo	RAMyo	LA	LV	Myo	RA	RV
Voxel2Mesh	94.1 $\pm$ 0.7	91.5 $\pm$ 1.7	83.6 $\pm$ 3.9	89.9 $\pm$ 3.2	89.3 $\pm$ 2.9	87.8 $\pm$ 2.7	86.4 $\pm$ 3.4	87.7 $\pm$ 2.0	87.3 $\pm$ 2.9
MeshDeformNet	93.9 $\pm$ 0.7	90.6 $\pm$ 1.2	83.4 $\pm$ 4.0	87.6 $\pm$ 4.1	89.2 $\pm$ 2.9	87.6 $\pm$ 3.0	86.1 $\pm$ 3.0	87.5 $\pm$ 1.8	86.9 $\pm$ 2.4
HeartDeformNet	93.9 $\pm$ 0.7	90.4 $\pm$ 1.4	83.1 $\pm$ 4.2	87.5 $\pm$ 4.2	89.1 $\pm$ 3.1	87.6 $\pm$ 3.0	86.1 $\pm$ 3.3	87.5 $\pm$ 2.1	87.2 $\pm$ 2.8
Ours 3p1	94.3 $\pm$ 1.1	92.1 $\pm$ 2.3	84.6 $\pm$ 4.6	90.1 $\pm$ 5.2	89.2 $\pm$ 2.8	88.2 $\pm$ 3.1	87.0 $\pm$ 3.3	87.5 $\pm$ 2.4	86.3 $\pm$ 3.3
Ours 3p6	95.3$\pm$0.7	94.2$\pm$0.9	87.7$\pm$2.5	92.2$\pm$3.8	89.5$\pm$2.6	88.7$\pm$2.8	87.6$\pm$3.0	88.3$\pm$2.3	87.3$\pm$3.3
Ablation-A	93.8 $\pm$ 0.8	90.5 $\pm$ 1.4	83.2 $\pm$ 4.2	87.6 $\pm$ 4.1	89.0 $\pm$ 3.3	87.3 $\pm$ 3.1	86.0 $\pm$ 3.3	87.3 $\pm$ 2.1	86.3 $\pm$ 3.1
Ablation-B	95.1 $\pm$ 0.7	93.4 $\pm$ 1.0	87.7 $\pm$ 1.4	91.9 $\pm$ 2.6	/	/	/	/	/
Ablation-C	93.1 $\pm$ 1.2	90.3 $\pm$ 1.5	78.5 $\pm$ 4.7	84.0 $\pm$ 4.8	87.9 $\pm$ 3.1	87.2 $\pm$ 2.9	86.8 $\pm$ 2.9	86.7 $\pm$ 2.9	84.7 $\pm$ 3.6

tetrahedral element-quality statistics (Sec. 3.1) to assess simulation readiness. To ensure a fair comparison, we plug all methods into the same warping module (our proposed warping pipeline). Across 400 test instances (each with ~ 3 M tetrahedra), ours HeartVolMesh yields 0.0% inverted elements, with a minimum scaled Jacobian of 0.0416 ± 0.0152 and a minimum dihedral angle of $5.56^\circ \pm 1.41^\circ$.

Surface meshes. Table 1 (right) reports surface-mesh accuracy before warping. Our Ours3p6 achieves the best CD/HD95 on all five surfaces and consistently strong NC, indicating improved localization under boundary ambiguity. The performance gap between Ours3p6 and Ours3p1 further supports the benefit of anisotropic covariance modelling for surface learning.

3.3 Ablation Study

Ablation-A replaces the covariance-guided NLL with Chamfer-only supervision. Ablation-B (same surface mesh result with Ours3p6) uses a coarser deformation-field for volumetric warping with reduced motion-coherence strength. Ablation-C removes mesh regularization (edge/Laplacian/normal) during surface learning. Ablation-C degrades both volumetric and surface performance, indicating that mesh regularization is important to avoid degenerate deformations that propagate to tetrahedral meshes. Ablation-A underperforms the full model, suggesting the benefit of covariance-guided supervision. Ablation-B shows mixed results, implying that overly coarse warping can oversmooth thin-wall regions.




Fig. 4. Quantitative summary on multi-centre in house data. Top: normalized radar plots over 24 volumetric parts. Bottom: boxplots over five surface structures. Ours3p6 shows lower errors and reduced variance.

4 Conclusion

We presented HeartVolMesh for cardiac volumetric mesh reconstruction with cross-case vertex correspondence from 3D images via covariance-guided graph deformation and template-conditioned warping. On multi-centre in-house data, our method consistently improves boundary accuracy. We trade end-to-end meshing (which typically yields a single fixed mesh specification after training) for a template-conditioned design that supports diverse simulation mesh specifications (e.g., density and element type). Future work will extend to cine CMR, benchmark multiple template specifications, and validate simulations.

Acknowledgements. AFF acknowledges support from the Royal Academy of Engineering under the RAEng Chair in Emerging Technologies (INSILEX CiET191919), ERC Advanced Grant–UKRI Frontier Research Guarantee (INSILICO EPY0304941), the UK Centre of Excellence on in-silico Regulatory Science and Innovation (UK CEiRSI) (10139527), the National Institute for Health and Care Research (NIHR) Manchester Biomedical Research Centre (BRC) (NIHR203308), the BHF Manchester Centre of Research Excellence (RE2430017), and the CRUK RadNet Manchester (C1994A28701). AZ/FL/AFF acknowledges support from NVIDIA Academic Grant Program.

Disclosure of Interests. AF is Co-founder and shareholder of OculomeX Health Ltd, adsilico Ltd, and Synaptive Consulting Ltd.

References

1. Viceconti, M., Henney, A., Morley-Fletcher, E.: In silico clinical trials: how computer simulation will transform the biomedical industry. *International Journal of Clinical Trials* **3**(2), 37–46 (2016). <https://doi.org/10.18203/2349-3259.ijct20161408>
2. Rodero, C., Baptiste, T.M.G., Barrows, R.K., et al.: A systematic review of cardiac in-silico clinical trials. *Progress in Biomedical Engineering* **5**(3), 032004 (2023). <https://doi.org/10.1088/2516-1091/acdc71>
3. Pathmanathan, P., Aycock, K., Badal, A., et al.: Credibility assessment of in silico clinical trials for medical devices. *PLoS Computational Biology* **20**(8), e1012289 (2024). <https://doi.org/10.1371/journal.pcbi.1012289>
4. Samei, E.: The future of in silico trials and digital twins in medicine. *PNAS Nexus* **4**(5), pgaf123 (2025). <https://doi.org/10.1093/pnasnexus/pgaf123>
5. Sel, K., Osman, D., Zare, F., et al.: Building Digital Twins for Cardiovascular Health: From Principles to Clinical Impact. *Journal of the American Heart Association* **13**(19), e031981 (2024). <https://doi.org/10.1161/JAHA.123.031981>
6. Updegrove, A., Wilson, N.M., Shadden, S.C.: SimVascular: An Open Source Pipeline for Cardiovascular Simulation. *Annals of Biomedical Engineering* **45**(3), 525–541 (2017). <https://doi.org/10.1007/s10439-016-1762-8>
7. Fedele, M., Quarteroni, A.: Polygonal surface processing and mesh generation tools for the numerical simulation of the cardiac function. *International Journal for Numerical Methods in Biomedical Engineering* **37**(4), e3435 (2021). <https://doi.org/10.1002/cnm.3435>
8. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3D surface construction algorithm. *Computer Graphics (Proc. SIGGRAPH)* **21**(4), 163–169 (1987). <https://doi.org/10.1145/37402.37422>
9. Updegrove, A., Wilson, N.M., Shadden, S.C.: Boolean and smoothing of discrete polygonal surfaces. *Advances in Engineering Software* **95**, 16–27 (2016). <https://doi.org/10.1016/j.advengsoft.2016.01.015>
10. Si, H.: TetGen, a Delaunay-based quality tetrahedral mesh generator. *ACM Transactions on Mathematical Software* **41**(2), Article 11, 1–36 (2015). <https://doi.org/10.1145/2629697>
11. Wickramasinghe, U., Remelli, E., Knott, G., Fua, P.: Voxel2Mesh: 3D mesh model generation from volumetric data. In: Martel, A.L., Abolmaesumi, P., Stoyanov, D., Mateus, D., Zuluaga, M.A., Zhou, S.K., Racocceanu, D., Joskowicz, L. (eds.) *MICCAI 2020*, LNCS, vol. 12264, pp. 299–308. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59719-1_30
12. Kong, F., Wilson, N.M., Shadden, S.C.: A deep-learning approach for direct whole-heart mesh reconstruction. *Medical Image Analysis* **74**, 102222 (2021). <https://doi.org/10.1016/j.media.2021.102222>
13. 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). <https://doi.org/10.1109/TMI.2022.3219284>
14. Gkioxari, G., Johnson, J., Malik, J.: Mesh R-CNN. In: 2019 IEEE/CVF International Conference on Computer Vision (ICCV), pp. 9784–9794. IEEE (2019). <https://doi.org/10.1109/ICCV.2019.00988>
15. Sorkine, O., Cohen-Or, D., Lipman, Y., Alexa, M., Rössl, C., Seidel, H.-P.: Laplacian Surface Editing. In: Scopigno, R., Zorin, D. (eds.) *Symposium on Geometry Processing*, pp. 179–188. The Eurographics Association (2004). <https://doi.org/10.2312/SGP/SGP04/179-188>

16. Wang, N., Zhang, Y., Li, Z., Fu, Y., Liu, W., Jiang, Y.-G.: Pixel2Mesh: Generating 3D Mesh Models from Single RGB Images. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) ECCV 2018, LNCS, vol. 11215, pp. 55–71. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01252-6_4
17. Kato, H., Ushiku, Y., Harada, T.: Neural 3D Mesh Renderer. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 3907–3916. IEEE (2018). <https://doi.org/10.1109/CVPR.2018.00411>
18. Lin, F., Xia, Y., Song, S., Ravikumar, N., Frangi, A.F.: High-throughput 3DRA segmentation of brain vasculature and aneurysms using deep learning. *Computer Methods and Programs in Biomedicine* **230**, 107355 (2023). <https://doi.org/10.1016/j.cmpb.2023.107355>
19. Hu, Y., Zhou, Q., Gao, X., Jacobson, A., Zorin, D., Panozzo, D.: Tetrahedral Meshing in the Wild. *ACM Transactions on Graphics* **37**(4), Article 60 (2018). <https://doi.org/10.1145/3197517.3201353>
20. Hu, Y., Schneider, T., Wang, B., Zorin, D., Panozzo, D.: Fast Tetrahedral Meshing in the Wild. *ACM Transactions on Graphics* **39**(4), Article 117 (2020). <https://doi.org/10.1145/3386569.3392385>
21. Kendall, A., Gal, Y., Cipolla, R.: Multi-Task Learning Using Uncertainty to Weigh Losses for Scene Geometry and Semantics. In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 7482–7491. IEEE (2018). <https://doi.org/10.1109/CVPR.2018.00781>
22. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation. In: Ourselin, S., Joskowicz, L., Sabuncu, M., Unal, G., Wells, W. (eds.) *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2016*, LNCS, vol. 9901, pp. 424–432. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46723-8_49
23. Umeyama, S.: Least-Squares Estimation of Transformation Parameters Between Two Point Patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **13**(4), 376–380 (1991). <https://doi.org/10.1109/34.88573>
24. Myronenko, A., Song, X.: Point Set Registration: Coherent Point Drift. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **32**(12), 2262–2275 (2010). <https://doi.org/10.1109/TPAMI.2010.46>
25. Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L.G., Leach, M.O., Hawkes, D.J.: Nonrigid Registration Using Free-Form Deformations: Application to Breast MR Images. *IEEE Transactions on Medical Imaging* **18**(8), 712–721 (1999). <https://doi.org/10.1109/42.796284>
26. Fan, H., Su, H., Guibas, L.J.: A Point Set Generation Network for 3D Object Reconstruction from a Single Image. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 2463–2471. IEEE (2017). <https://doi.org/10.1109/CVPR.2017.264>
27. Dou, H., Virtanen, S., Ravikumar, N., Frangi, A.F.: A Generative Shape Compositional Framework to Synthesize Populations of Virtual Chimeras. *IEEE Transactions on Neural Networks and Learning Systems* **36**(3), 4750–4764 (2025). <https://doi.org/10.1109/TNNLS.2024.3374121>
28. Lin, F., Zakeri, A., Xue, Y., MacRaid, M., Dou, H., Zhou, Z., Zou, Z., Sarrami-Faroushani, A., Duan, J., Frangi, A.F.: From Pixels to Polygons: A Survey of Deep Learning Approaches for Medical Image-to-Mesh Reconstruction. *Medical Image Analysis* **113**, 104153 (2026). <https://doi.org/10.1016/j.media.2026.104153>