



3D Mitral Valve Surface Reconstruction from 3D TEE via Graph Neural Networks

Matthias Ivantsits¹(✉), Boris Pfahringer¹, Markus Huellebrand^{1,2}, Lars Walczak^{1,2}, Lennart Tautz², Olena Nemchyna³, Serdar Akansel³, Jörg Kempfert³, Simon Sündermann¹, and Anja Hennemuth^{1,2,4}

¹ Charité – Universitätsmedizin Berlin, Augustenburger Pl. 1, 13353 Berlin, Germany
matthias.ivantsits@charite.de

² Fraunhofer MEVIS, Am Fallturm 1, 28359 Bremen, Germany

³ German Heart Center Berlin, Augustenburger Pl. 1, 13353 Berlin, Germany

⁴ DZHK (German Centre for Cardiovascular Research), Berlin, Germany

Abstract. Mitral valve insufficiency is a condition in which the valve does not close properly, and blood leaks back into the atrium from the ventricle. Valve assessment for surgery planning is typically performed with 3D transesophageal echocardiography (TEE). The simulation of the resulting valve dynamics can support selecting the most promising surgery strategy. These simulations require an accurate reconstruction of the open valve as a 3D surface model. 3D mitral valve reconstruction from image data is challenging due to the fast-moving and thin valve leaflets, which might appear blurred and covered by very few voxels depending and spatio-temporal resolution. State-of-the-art voxel-based CNN segmentation methods need an additional processing step to reconstruct a 3D surface from this voxel-based representation which can introduce unwanted artifacts. We propose an end-to-end deep-learning-based method to reconstruct a 3D surface model of the mitral valve directly from 3D TEE images. The suggested method consists of a CNN-based voxel encoder and decoder inter-weaved with a graph neural network-based (GNN) multi-resolution mesh decoder. This GNN samples feature vectors from the CNN-decoder at different resolutions to deform a prototype mesh. The model was trained on 80 sparsely annotated 3D TEE images (1 mm³ voxel resolution) of the valve during end-diastole. Each time frame was annotated by two cardiovascular experts on nine planes rotated around the axis through the apex of the left ventricle and the center of the mitral valve. Our method's average bidirectional point-to-point distance is 1.1 mm, outperforming the inter-observer point-to-point distance of 1.8 mm.

Keywords: Mitral valve reconstruction · Deep learning · Graph neural network

1 Introduction

The mitral valve regulates blood flow between the left atrium and left ventricle. It consists of the anterior and posterior leaflet, separated by two indentations—the

anterior and posterior commissure (Fig. 1A). These two leaflets are attached to the annulus and are pulled down during the diastole via thin chords attached to the papillary muscles, which are connected to the left myocardium. During systole, the myocardium contracts, and the mitral valve closes (Fig. 1A). Mitral valve regurgitation is the most common mitral valve disease—with approximately 2% affected people worldwide [1]—causing blood flow from the ventricle into the atrium during systole. This regurgitation can be classified into three main types—type I, II, III—defined by Carpentier [2, 3]. These mitral valve regurgitation types need different methods for repair, or replacement [4]. A common method to assess the pathology of the mitral valve is via 3D transesophageal echocardiography (TEE) [5, 6]. Commercially available 3D TEE mitral valve tools typically analyze the mitral valve pathology or morphology in a closed state [7–9]. This allows the calculation of the annulus perimeter, annulus area, anterior and posterior leaflet area and length, and tenting height [10] (Fig. 1B). Assessing the mitral valve leaflets in a closed state is challenging due to their connection in the coaptation area. To enable the assessment of the valve leaflets and the characteristics of the closed valve, recent approaches reconstruct the valve model in an open state and propagate the valve model to the closed state [11, 12]. The model can be changed according to different treatment strategies to select the most promising surgery procedure. This approach requires an initial reconstruction of the mitral valve from image data as a 3D surface. This reconstruction is challenging because the fast-moving and thin valve leaflets might be covered by only 1–2 voxels and blurred depending on the spatiotemporal sampling. State-of-the-art voxel-based CNN approaches [13–16] or interactive methods to delineate the valve [17, 18] need an additional processing step to reconstruct a 3D surface from this voxel-based representation, which can introduce unwanted artifacts. We propose an end-to-end deep-learning-based method to reconstruct a 3D surface model of the open mitral valve directly from 3D TEE images.

2 Materials and Method

2.1 Dataset

The dataset was acquired utilizing the GE Vingmed Ultrasound Vivid E9 system in the German Heart Centre Berlin (DHZB). The dataset consists of 3D TEE images of 38 patients examined before and after surgery. 81 end-diastolic image volumes were used for model training and validation. The temporal resolution is between 12 and 65 ms in-plane resolution between 0.47 and 1.43 mm, and through-plane resolution between 0.33 and 0.98 mm. The image volumes were sparsely annotated by two cardiologists with four and 17 years of experience as a cardiologist, utilizing the method proposed by Tautz et al. [19]. The leaflets were annotated on nine planes rotated around the axis through the apex of the left ventricle and the center of the mitral valve (Fig. 1B). Additionally, the annulus and the orifice of the mitral valve were annotated. From these leaflet poly-lines,

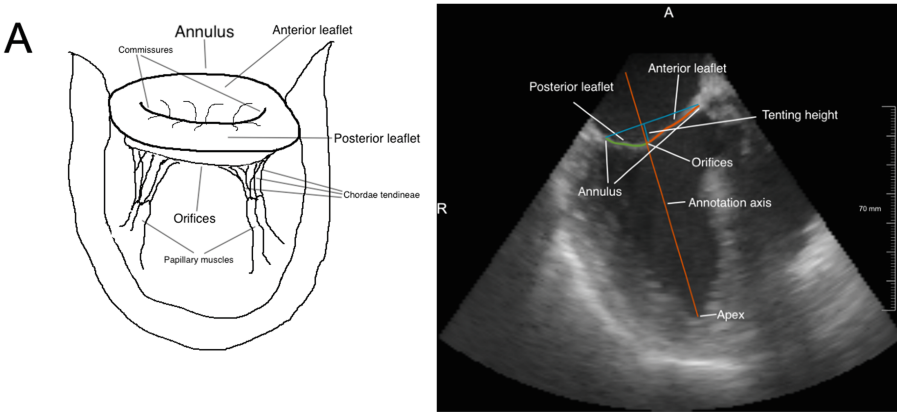


Fig. 1. An illustration of the relevant anatomical structures of the mitral valve and supporting structures (A). B illustrates relevant anatomical structures plus clinically important parameters.

a full 3D surface mesh was reconstructed for the final valve geometry, with each mesh consisting of 280 to 1200 nodes.

2.2 Method

Our proposed method is derived from the Voxel2Mesh [20] architecture, which is an extension to Pixel2Mesh [21] and Pixel2Mesh++ [22]. This method takes a 3D image in the voxel domain as input and produces the valve as a 3D surface model without any post-processing steps. The model consists of three main blocks—a voxel encoder and decoder and a mesh decoder—which are illustrated in Fig. 2. The voxel encoder and decoder are basic blocks, as proposed in the UNet paper [23], which compress the data into a latent space and reverse the process. The mesh decoder takes a prototype mesh as additional input, which is successively deformed. During each step of the mesh decoder, features around a fixed sphere of the mesh node positions are sampled and aggregated by 1D convolution. These latent node features plus the actual mesh node positions are concatenated and passed to a graph neural network [24], which produces a 3D offset vector for each mesh node. This delta is added to the previous mesh to produce an intermediate, low-resolution mesh. Next, this mesh is up-sampled by adding mesh nodes in each edge’s center, creating four new faces. This process is repeated until the full resolution of the voxel representation is reached.

We extended the Voxel2Mesh architecture by adapting the mesh decoder to deform a topological annulus (we are referring to a topological/geometrical object here) instead of a sphere, which is, aside from a torus, the only feasible topology for a mitral valve surface reconstruction, due to the preservation of the prototype topology. The proposed method uses a uniform up-sampling strategy instead of the adaptive unpooling method suggested in the Voxel2Mesh

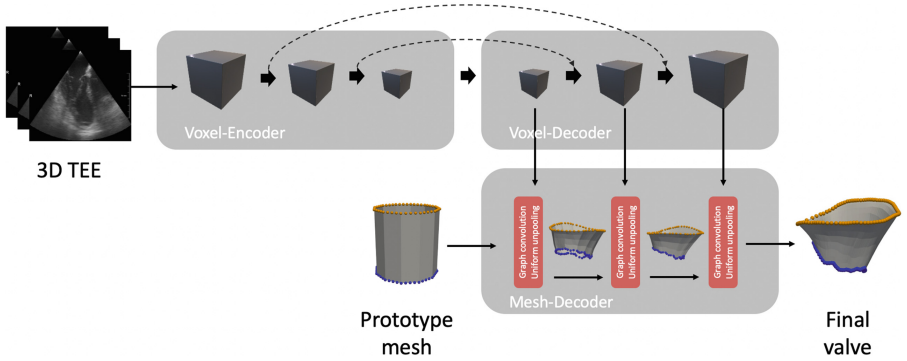


Fig. 2. An overview of the proposed architecture, including the three main blocks—a voxel-encoder, a voxel-decoder, and a mesh-decoder. The model takes a 3D TEE image as input and produces a 3D surface model of the mitral valve. The mesh-decoder samples features from the latent voxel-space and successively refines the prototype mesh at increasing resolutions. During the deformation, the orientation of the edge vertices is preserved, thus reconstructing the mitral valve with known annulus and orifice vertices. (Color figure online)

architecture [20], which does not preserve the topology of an annulus. Furthermore, we implement a dedicated loss function for mitral valve reconstruction. This loss function ensures an unfolded 3D surface and guarantees a valve model with labeled nodes, where the annulus and the orifices of the final mesh are known. This assignment of anatomical structures is illustrated in Fig. 2, where the annulus vertices (yellow) and orifice vertices (blue) are known in the prototype mesh, as well as in the final reconstructed mitral valve surface, resulting in a known orientation.

The final loss function L is a linear combination of multiple loss terms γ weighted with their respective hyperparameters λ at different resolutions R plus the binary cross-entropy loss (illustrated in Eq. 1.6). The set of loss terms Γ includes the chamfer loss γ_{MV} , defined as the average minimum distance between two point clouds. Furthermore, the same loss function is applied to the annulus and orifice vertices to ensure the correct orientation of the final 3D surface. Additionally, the loss contains constraints for the edge length γ_E , normal consistency of neighboring faces γ_{NC} , Laplacian smoothing objective γ_L , and a face normal loss γ_{MVN} , which ensures similar normal vectors of the predicted and ground-truth faces. These additional constraints ensure a smooth and non-self-folding solution to the 3D surface reconstruction.

$$\gamma_{MV} = \sum_{v_p \in S_p} \min_{v_g \in S_g} \|v_p - v_g\|_2^2 + \sum_{v_p \in S_g} \min_{v_g \in S_p} \|v_p - v_g\|_2^2 \quad (1)$$

Equation 1: The chamfer loss function used for the mitral valve, the annulus, and the orifices. S_p and S_g denote the predicted and ground-truth mitral valve surface models, where v_p and v_g designate vertices sampled from the surface.

$$\gamma_E = \frac{1}{|E|} \sum_{(v_0, v_1) \in E} (\|v_0 - v_1\|_2 - L_t)^2 \quad (2)$$

Equation 2: The edge length loss, where E denotes the set of edges of a mesh. v_0 and v_1 are the vertices composing the edge, and L_t is defined as a constant target length.

$$\gamma_{NC} = \frac{1}{|N_F|} \sum_{(n_0, n_1) \in N_F} 1 - \frac{n_0 \cdot n_1}{\|n_0\| \cdot \|n_1\|} \quad (3)$$

Equation 3: The normal consistency loss, where N_F denotes the set of neighboring faces of the predicted mesh. n_0 and n_1 designate the normal vectors of these neighboring faces.

$$\gamma_L = \frac{1}{|V|} \sum_{v \in V} \left\| \sum_{v' \in N(v)} v - v' \right\| \quad (4)$$

Equation 4: The Laplacian smoothing constraint, where V denotes the set of vertices of the predicted mitral valve surface. $N(v)$ depicts the set of neighboring vertices of vertex v .

$$\gamma_{MVN} = \frac{1}{|F_p|} \sum_{f_p \in F_p} 1 - \frac{n_p \cdot n_g}{\|n_p\| \cdot \|n_g\|} \quad (5)$$

Equation 5: The face normal loss, where F_p depicts the set of faces of the predicted mitral valve. n_p denotes the normal vector of face f_p and n_g the normal vector of the closest face f_g to the predicted face f_p .

$$L = \lambda_{CE} \gamma_{CE} + \sum_{\gamma \in \Gamma} \sum_{r \in R} \lambda_{\gamma r} \gamma_r \quad (6)$$

Equation 6: The loss function of the proposed method, including terms for the cross-entropy loss γ_{CE} , the chamfer loss of the mitral valve γ_{MV} , the edge length loss γ_E , the normal consistency loss γ_{NC} , the Laplacian smoothing objective γ_L , and a face normal loss γ_{MVN} . Additionally, each loss term is weighted by a hyperparameter λ .

3 Results

We performed all experiments on an Intel(R) Core(TM) i7-8700K CPU @ 3.70 GHz with 16 GBs RAM and an Nvidia RTX 2080 Ti GPU with 11 GB memory. 63 3D TEE volumes were used for model training, and the remaining 18 for model testing. We performed a 5-fold cross-validation (CV) to find optimal hyperparameters for the number of down- and up-sampling steps R and the loss weights λ . Before model training and testing, all images were re-sampled to an isotropic voxel-spacing of 1 mm^3 . We performed random affine, blur, and Gaussian noise augmentation during model training. Furthermore, a grid search on

the loss hyperparameters was performed, combined with visual examinations, to avoid folded surface reconstructions. The final loss hyperparameters were uniformly weighted across all resolutions and weighted with: $\lambda_{MV} = 1$, $\lambda_A = 1$, $\lambda_O = 1$, $\lambda_{MVN} = 0.2$, $\lambda_L = 0.2$, $\lambda_{NC} = 0.2$, $\lambda_E = 0.2$, and $\lambda_{CE} = 1$.

We performed an interobserver variability analysis for the 81 3D volumes annotated by both experts. Figure 3 illustrates the average and 95th-percentile bidirectional point-to-point distance of the annotations to the reconstructed mitral valve mesh. We observe an average distance of 1.1 mm and 95th-percentile of 2.13 mm outperforming the inter-observer metrics of 1.86 mm and 3.82 mm. This figure highlights the average and 95th-percentile distance of the annulus and orifices, with an average distance of 1.29 mm and an average distance of the orifices of 1.67 mm, compared to the inter-observer metrics of 2.45 mm and 3.29 mm. The 95th-percentile results in 2.23 mm for the annulus and 3.19 mm for the orifices compared to the experts' variability of 4.06 mm and 6.71 mm.




Fig. 3. An illustration of the bidirectional inter-observer and surface reconstruction metrics. This comparison includes the average and 95th-percentile point-to-point distance of the reconstructed valve to the ground truth and the distances for the annulus and orifices.

Additionally, we performed a linear regression on the target and predicted mitral valve area plus the maximum annulus diameter to identify any bias in the proposed model. We observe an R^2 value of 0.91 for the mitral valve area and 0.88 for the annulus diameter. The average difference in the mitral valve area




Fig. 4. An illustration of the linear regression of the mitral valve surface area and the maximum annulus diameter. The linear regression of the surface area results in an R^2 value of 0.91 and a bias of -103 mm^2 . The linear regression of the maximum annulus diameter results in an R^2 value of 0.88 and a bias of -1.9 mm .




Fig. 5. An illustration of two 3D surface mitral valve reconstructions produced by our model. The orange contours represent the reconstructed mitral valves produced by the proposed method, and the red contours the reconstruction created by the cardiologists. Row A depicts a result with an average point-to-point distance of 0.37 mm. Row B illustrates a reconstruction with an average point-to-point distance of 1.23 mm. (Color figure online)

is -103 mm^2 . The average difference of the maximum annulus diameter results -1.9 mm .

Figure 5 illustrates two 3D surface reconstruction results of the proposed method. Row A illustrates a mitral valve reconstruction that aligns very well with the reconstruction by the cardiologist. The average bidirectional point-to-point distance of this result is 0.37 mm. The sagittal, coronal, and axial planes show almost perfect alignment. Row B depicts an example with an average point-to-point distance of 1.23 mm, where the sagittal and coronal planes reveal large

deviations towards the orifices. The axial plane illustrates a large discrepancy in the annulus.

4 Discussion and Conclusion

We have presented an end-to-end deep-learning method to reconstruct a 3D surface model of the mitral valve from 3D TEE images without post-processing. By directly reconstructing the 3D surface, we avoid artifacts introduced in conventional approaches. The inter-observer study performed and illustrated in Fig. 3 exhibits an average bidirectional point-to-point distance of 1.1 mm, surpassing the average distance by two cardiologists of 1.86 mm. The analysis shown in Fig. 4 reveals a bias with an under-estimation of the reconstructed surface area of -103 mm^2 and -1.9 mm of the maximum annulus diameter. Furthermore, Fig. 5 B illustrates a reconstruction with an average point-to-point distance of 1.23 mm. The sagittal and coronal views reveal large discrepancies towards the orifices of the annulus.

To improve the proposed method, it is very likely that an increase in sample size can help better reconstruct the 3D surface model of the mitral valve. The inter-observer variability analysis highlights that an accurate annotation of the mitral valve from 3D TEE images is complex, especially the distinction between the chordae tendineae and the leaflets. Since acquiring ground truths for the valve is difficult, an extension of the training set with synthetic data might be advantageous. Future extensions to this method might also include a vertex classification of the reconstructed 3D surface. This extension will lead to an extended semi-registered model, where the individual segments of the leaflets are learned end-to-end.

References

1. Douedi, S., Douedi, H.: Mitral regurgitation. In: StatPearls. StatPearls Publishing, Treasure Island (FL) (2022). <https://www.ncbi.nlm.nih.gov/books/NBK553135/>. Accessed 8 Nov 2021
2. Apostolidou, E., Maslow, A.D., Poppas, A.: Primary mitral valve regurgitation: update and review. *Glob. Cardiol. Sci. Pract.* **2017**(1), e201703 (2017). <https://doi.org/10.21542/gcsp.2017.3>
3. Lancellotti, P., Moura, L., Pierard, L.A., et al.: European association of echocardiography recommendations for the assessment of valvular regurgitation. Part 2: mitral and tricuspid regurgitation (native valve disease). *Eur. J. Echocardiogr.* **11**(4), 307–332 (2010). <https://doi.org/10.1093/ejechocard/jeq031>
4. Allen, N., O’Sullivan, K., Jones, J.M.: The most influential papers in mitral valve surgery; a bibliometric analysis. *J. Cardiothorac. Surg.* **15**, 175 (2020). <https://doi.org/10.1186/s13019-020-01214-y>
5. Nishimura, R.A., Otto, C.M., Bonow, R.O., et al.: 2017 AHA/ACC focused update of the 2014 AHA/ACC guideline for the management of patients with valvular heart disease: a report of the American college of cardiology/American heart association task force on clinical practice guidelines. *Circulation* **135**(25), e1159–e1195 (2017). <https://doi.org/10.1161/CIR.0000000000000503>

6. de Groot-de Laat, L.E., McGhie, J., Ren, B., Frowijn, R., Oei, F.B., Geleijnse, M.L.: A modified echocardiographic classification of mitral valve regurgitation mechanism: the role of three-dimensional echocardiography. *J. Cardiovasc. Imaging* **27**(3), 187–199 (2019). <https://doi.org/10.4250/jcvi.2019.27.e29>
7. Sandra, L., et al.: A new diagnostic tool for effective regurgitant orifice quantification in mitral regurgitation. *Echocardiography* **35**(11), 1812–1817 (2018). <https://doi.org/10.1111/echo.14114>
8. Veronesi, F., Lie, G., Rabben, S.: 4D auto MVQ. GE Healthcare (2017)
9. Tomtec. 4D MV-Assessment (2018)
10. Ryan, L., et al.: Quantification and localization of mitral valve tenting in ischemic mitral regurgitation using real-time three-dimensional echocardiography. *Eur. J. Cardiothorac. Surg.* **31**(5), 839–844 (2007). <https://doi.org/10.1016/j.ejcts.2007.01.050>
11. Walczak, L., et al.: Using position-based dynamics for simulating mitral valve closure and repair procedures. *Comput. Graph. Forum* **41**, 270–287 (2022). <https://doi.org/10.1111/cgf.14434>
12. Walczak, L., et al.: Interactive editing of virtual chordae tendineae for the simulation of the mitral valve in a decision support system. *Int. J. Comput. Assist. Radiol. Surg.* **16**(1), 125–132 (2020). <https://doi.org/10.1007/s11548-020-02230-y>
13. Carnahan, P., Moore, J., Bainbridge, D., Eskandari, M., Chen, E.C.S., Peters, T.M.: DeepMitral: fully automatic 3D echocardiography segmentation for patient specific mitral valve modelling. In: de Bruijne, M., et al. (eds.) MICCAI 2021. LNCS, vol. 12905, pp. 459–468. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-87240-3_44
14. Costa, E., et al.: Mitral valve leaflets segmentation in echocardiography using convolutional neural networks. In: 2019 IEEE 6th Portuguese Meeting on Bioengineering (ENBENG), pp. 1–4 (2019). <https://doi.org/10.1109/ENBENG.2019.8692573>
15. Kerfoot, E., Clough, J., Oksuz, I., Lee, J., King, A.P., Schnabel, J.A.: Left-ventricle quantification using residual U-net. In: Pop, M., et al. (eds.) STACOM 2018. LNCS, vol. 11395, pp. 371–380. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-12029-0_40
16. Liu, S., et al.: 3D anisotropic hybrid network: transferring convolutional features from 2D images to 3D anisotropic volumes. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) MICCAI 2018. LNCS, vol. 11071, pp. 851–858. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00934-2_94
17. Carnahan, P., et al.: Interactive-automatic segmentation and modelling of the mitral valve. In: Coudière, Y., Ozenne, V., Vigmond, E., Zemzemi, N. (eds.) FIMH 2019. LNCS, vol. 11504, pp. 397–404. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-21949-9_43
18. Jassar, A.S., Brinster, C.J., Vergnat, M., et al.: Quantitative mitral valve modeling using real-time three-dimensional echocardiography: technique and repeatability. *Ann. Thorac. Surg.* **91**(1), 165–171 (2011). <https://doi.org/10.1016/j.athoracsur.2010.10.034>
19. Tautz, L., et al.: Extraction of open-state mitral valve geometry from CT volumes. *Int. J. Comput. Assist. Radiol. Surg.* **13**(11), 1741–1754 (2018). <https://doi.org/10.1007/s11548-018-1831-6>
20. Wickramasinghe, U., Remelli, E., Knott, G., Fua, P.: Vox2Mesh: 3D mesh model generation from volumetric data. In: Martel, A.L., et al. (eds.) MICCAI 2020. LNCS, vol. 12264, pp. 299–308. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59719-1_30

21. Wang N., Zhang Y., Li Z., Fu Y., Liu W., Jiang Y.: Pixel2Mesh: generating 3D mesh models from single RGB images. In: European Conference on Computer Vision (2018)
22. Wen C., Zhang Y., Li Z., Fu Y.: Pixel2Mesh++: multi-view 3D mesh generation via deformation. In: International Conference on Computer Vision (2019)
23. Ronneberger, O., Fischer, P., Brox, T.: U-net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
24. Kipf, T., Welling, M.: Semi-supervised classification with graph convolutional networks. In: 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, 24–26 April 2017, Conference Track Proceedings (2017)