

Reconstructing 3D Cardiac Anatomies from Misaligned Multi-View Magnetic Resonance Images with Mesh Deformation U-Nets

Marcel Beetz

MARCEL.BEETZ@ENG.OX.AC.UK

Abhirup Banerjee

ABHIRUP.BANERJEE@ENG.OX.AC.UK

Vicente Grau

VICENTE.GRAU@ENG.OX.AC.UK

*Institute of Biomedical Engineering
Department of Engineering Science
University of Oxford
Oxford, OX3 7DQ, UK*

Editor: —

Abstract

High-quality three-dimensional (3D) representations of cardiac anatomy and function are crucial for improving cardiac disease diagnosis beyond strictly volume-based biomarkers used in current clinical practice, as well as for the accurate simulation of cardiac electrophysiology and mechanics. However, current gold standard cardiac magnetic resonance imaging (MRI) protocols typically only acquire a set of 2D slices to approximate the true 3D morphology of the underlying heart. In this work, we propose a novel geometric deep learning method, the *Mesh Deformation U-Net*, to reconstruct 3D cardiac surface meshes from 2D MRI slices as the key part of a fully automatic end-to-end pipeline. Its architecture combines spectral graph convolutions and mesh sampling operations in a hierarchical encoder-decoder structure to enable efficient multi-scale feature learning directly on mesh data. A targeted preprocessing step approximately fits a template mesh to the sparse MRI contours, before the Mesh Deformation U-Net corrects for motion-induced slice misalignment by simultaneously utilising information from multiple MRI views and the template-induced anatomical shape prior. We evaluate the Mesh Deformation U-Net on a large synthetic dataset of heart anatomies and outperform multiple benchmark approaches while achieving small reconstruction errors below the pixel size of the underlying image resolution for three different cardiac substructures. Furthermore, we apply the pre-trained Mesh Deformation U-Net as the key component of a 4-step reconstruction pipeline to cine magnetic resonance images of the UK Biobank and observe realistic heart reconstructions on both a local and global level. We calculate multiple widely used clinical metrics for the reconstructed meshes and obtain values in line with other large-scale population studies.

Keywords: Cardiac Surface Reconstruction, Mesh Deformation U-Nets, Cine MRI, Slice Misalignment Correction, Motion Artifact, Spectral Graph Convolutions, Mesh Pooling, Geometric Deep Learning

1. Introduction

Cardiac magnetic resonance imaging (MRI) is one of the gold standard medical imaging modalities for the diagnosis of cardiac anatomy and function (Stokes and Roberts-Thomson,

2017). The current clinical practice, mostly the cardiac cine MRI protocol, usually images the heart myocardium with high soft-tissue contrast in a collection of sparse and intersecting (a stack of short-axis and few long-axis) 2D image planes, with a typical $1\text{-}2 \times 1\text{-}2 \text{ mm}^2$ in-plane resolution and 8-10 mm out-of-plane resolution. However, this can only present the underlying heart anatomy at a finite number of spatial locations and orientations and as a result, provides a sparse representation of the true 3D geometry of the human heart, limiting the accuracy of cardiovascular disease diagnosis (Corral Acero et al., 2022; Di Folco et al., 2022; Suinesiaputra et al., 2017).

In order to solve the sparsity issue of 2D cine MRI, many previous works have attempted to reconstruct the complete 3D cardiac surfaces from the acquired 2D slices (Banerjee et al., 2021a, 2022; Joyce et al., 2022; Mauger et al., 2018, 2019; Villard et al., 2018). One of the biggest challenges for 3D cardiac surface reconstruction from sparse 2D cine MRI acquisitions is the presence of slice misalignment artifacts induced by the respiratory or breathing motion, cardiac or heart-beating motion, and the movements caused by the patient or imaging device. Several previous approaches have aimed to first segment the acquired 2D images in the cardiac substructures of interest and then address the misalignment artifacts by optimising the consistency among the sparse heart contours, thus accounting mainly for the in-plane misalignment artifacts (McLeish et al., 2002; Su et al., 2014; Villard et al., 2017). A recent approach aims to fit a statistical shape model (SSM) (Bai et al., 2015) to the sparse heart contours in 3D space and then optimally align the contours on the SSM, accounting for both in-plane and out-of-plane misalignments (Banerjee et al., 2021b). The final 3D reconstruction usually relies on fitting a smooth surface mesh to the resulting contours in a per-case regularised optimisation procedure (Banerjee et al., 2021a; Joyce et al., 2022; Lamata et al., 2014; Mauger et al., 2019, 2018; Villard et al., 2018).

Although these methods can accurately model the patient-specific 3D deformations for cardiac surface modeling, they suffer from extended execution times, making them unsuitable for population-level analysis. Recently, deep learning based algorithms have become the state-of-the-art approaches for 3D cardiac surface reconstruction tasks, primarily due to their faster execution times, easy scalability, high versatility, and good performance based on the pretrained models. However, these approaches have not been applied to real data (Xu et al., 2019), can only process individual input images (Wang et al., 2020; Zhou et al., 2019), depend on inefficient representations of anatomical surface data in the form of voxel-grids (Chen et al., 2021; Xu et al., 2019), have been validated on only a small number of real cases (Beetz et al., 2021a), or require additional complex preprocessing and postprocessing steps that are prone to errors (Beetz et al., 2021a; Chen et al., 2021).

In this work, we propose the *Mesh Deformation U-Net* as a novel geometric deep learning approach to cardiac surface reconstruction. Its architecture follows a hierarchical encoder-decoder design with U-Net-inspired (Ronneberger et al., 2015) long skip connections and is specifically designed to directly process anatomical surface mesh data in an efficient manner. This enables the network to overcome the considerable memory and execution inefficiencies of previous voxelgrid-based deep learning (Xu et al., 2019) and per-case optimisation approaches (Banerjee et al., 2021a; Joyce et al., 2022; Mauger et al., 2018; Villard et al., 2018) respectively. In addition, a targeted template mesh fitting preprocessing step introduces an anatomical shape prior in a fast and straightforward manner which is particularly beneficial for the highly sparse anatomical input data. Furthermore, it allows the Mesh Deformation

U-Net to retain the same vertex connectivity for all meshes in the dataset which is an important requirement for many follow-up 3D cardiac modeling tasks (Beetz et al., 2021b,c, 2022a,b; Corral Acero et al., 2022; Li et al., 2022; Mauger et al., 2019).

2. Methods

2.1 Dataset and Preprocessing

We use two datasets in this work for method development and evaluation. The first one is a synthetic dataset derived from a 3D MRI-based statistical shape model (SSM) (Bai et al., 2015). Hereby, we first generate a population of 250 3D cardiac anatomies from the SSM which we consider as our ground truth 3D shapes for the reconstruction task, since they are based on high-resolution 3D MRI acquisitions ($1.25 \times 1.25 \times 2$ mm). Next, we slice each of the meshes at the common imaging planes of a typical cine MRI acquisition and then apply 10 different random rigid transformations to each slice to introduce slice misalignment that mimics real acquisitions, resulting in sparse and misaligned point cloud representations of the underlying anatomy. Since the Mesh Deformation U-Net requires 3D triangular meshes with the same vertex connectivity as its inputs, we design an approximate mesh fitting step to convert the sparse anatomy point clouds into dense meshes. To this end, we select the mean mesh of the SSM as our template mesh and deform it to fit the respective sparse point clouds. We determine the deformation based on the Earth Mover’s distance (EMD) between the sparse point cloud and the template mesh. The EMD provides a one-to-one mapping between points and vertices by minimising the average mapping distance between point cloud P_1 and mesh vertices P_2 , as follows:

$$EMD(P_1, P_2) = \min_{\phi: P_1 \rightarrow P_2} \frac{1}{|P_1|} \sum_{x \in P_1} \|x - \phi(x)\|_2. \quad (1)$$

We move each vertex of the template mesh to the spatial location of the corresponding closest point of the point cloud as determined by the Earth Mover’s distance in our deformation step. This results in an approximate fit of the template mesh to the sparse contours which includes all anatomical information of the MRI segmentation in the initial fitted mesh, while maintaining the template-induced built-in shape prior in a dense representation of the cardiac surface.

The second dataset consists of 1000 real cine MR images of the UK Biobank study (Petersen et al., 2013) with a pixel resolution of $1.8 \times 1.8 \times 8.0$ mm. In order to apply the Mesh Deformation U-Net to this real dataset, we develop the 4-step reconstruction pipeline described in Sec. 2.4.

2.2 Mesh Deformation U-Net

The architecture of the proposed Mesh Deformation U-Net utilises recent advances in geometric deep learning in a hierarchical encoder-decoder structure (Fig. 1).

Its inputs and outputs are 3D surface meshes with the same vertex-connectivity, and the vertices are stored as $n \times 3$ vectors where n refers to the number of vertices and 3 to the x,y,z coordinates of each vertex respectively. In our experiments, we apply the network to three cardiac substructures and use a different number of vertices, namely 1870, 2620, and 1706,

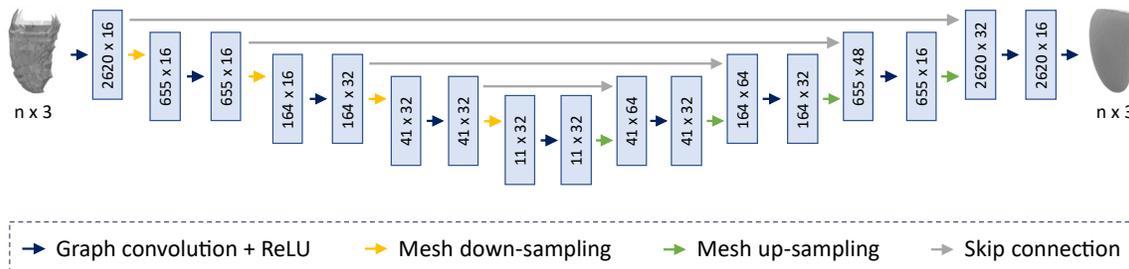


Figure 1: Architecture of the Mesh Deformation U-Net. It follows a hierarchical encoder-decoder structure with U-Net-inspired long skip connections and design similarities between corresponding levels. Graph convolutions and mesh sampling operations enable multi-scale feature learning directly on mesh data.

for the LV endocardial, LV epicardial, and RV endocardial anatomy meshes, respectively. The network architecture consists of multiple blocks of graph convolution, rectified linear unit (ReLU) activation layers, and mesh sampling operations organised in a hierarchical structure to enable efficient multi-scale feature learning directly on mesh data. Similar to the U-Net (Ronneberger et al., 2015), the encoder and decoder follow a symmetric design, where corresponding levels have the same mesh resolution and number of feature maps and are connected by long skip connections to facilitate the information flow between earlier and later parts of the network. We use spectral graph convolutions for feature learning which we calculate via the Chebyshev polynomial approximation of order 5 (Defferrard et al., 2016), while the sampling operations rely on quadric error minimisation (Ranjan et al., 2018).

2.3 Network Training and Implementation

We select the vertex-wise mean squared error between the predicted meshes and ground truth meshes as the loss function for the Mesh Deformation U-Net. We choose the Adam optimiser (Kingma and Ba, 2014) with a learning rate of 0.001 and a batch size of 8 to train all networks on a CPU with 8 GB memory. We employ a dataset split of 75%/5%/20% for the training, validation, and testing dataset, respectively, and stop the training process when no loss improvement has been observed on the validation dataset for 10 epochs. We use the PyTorch (Paszke et al., 2019) and PyTorch Geometric (Fey and Lenssen, 2019) frameworks to implement our deep learning and geometric deep learning code, respectively.

2.4 Cardiac Anatomy Reconstruction Pipeline

In order to utilise the Mesh Deformation U-Net for cardiac anatomy reconstruction from real cardiac MRI acquisitions, we propose the 4-step pipeline depicted in Fig. 2.

In the first step, we segment the input MR images into four classes (LV endocardium and epicardium, RV endocardium, background) using three separate convolutional neural networks for the short-axis (SAX), 4 chamber long-axis (LAX), and 2 chamber LAX views, respectively (Bai et al., 2018). Next, we extract the contours of the segmentation masks for each cardiac substructure and place the combined information from each view in 3D

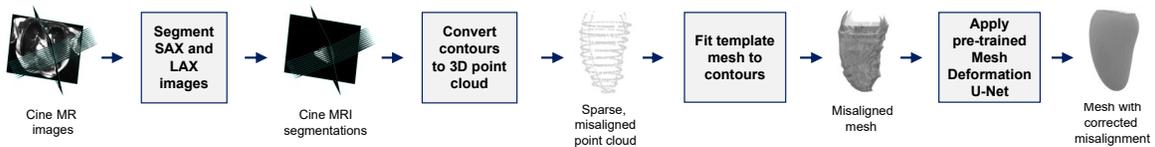


Figure 2: Overview of the proposed 4-step cardiac anatomy reconstruction pipeline. The raw cine MRI acquisitions are first segmented by convolutional neural networks, and the resulting cardiac anatomy contours are placed in 3D space. Then, a template mesh is approximately fitted to the 3D contours before the Mesh Deformation U-Net corrects the slice misalignment in the key step of the pipeline.

space as a point cloud (Banerjee et al., 2021a). Then, we fit the template mesh of the SSM dataset to this sparse point cloud representation using the mesh fitting procedure described in Sec. 2.1 to obtain a dense anatomy mesh. Finally, a Mesh Deformation U-Net is applied to remove any slice misalignment from the mesh in the key step of the pipeline and output the final reconstructed 3D cardiac shape.

3. Experiments

3.1 Reconstruction from Synthetic MRI Contours

First, we want to assess whether the Mesh Deformation U-Net combined with the approximate mesh fitting step is capable of accurately reconstructing synthetic 3D cardiac shapes from sparse and misaligned contours. To this end, we first apply the mesh fitting step to all sparse anatomy point clouds of the SSM dataset (Sec. 2.1). Next, we train a Mesh Deformation U-Net using the fitted meshes and the ground truth meshes of the training dataset. Then, we select the fitted meshes of the test dataset and pass them through the trained Mesh Deformation U-Net to obtain its predicted mesh reconstructions on unseen data. We execute this process separately for the LV endocardium and epicardium and RV endocardium, and present the results of three sample cases for each substructure in Fig. 3.

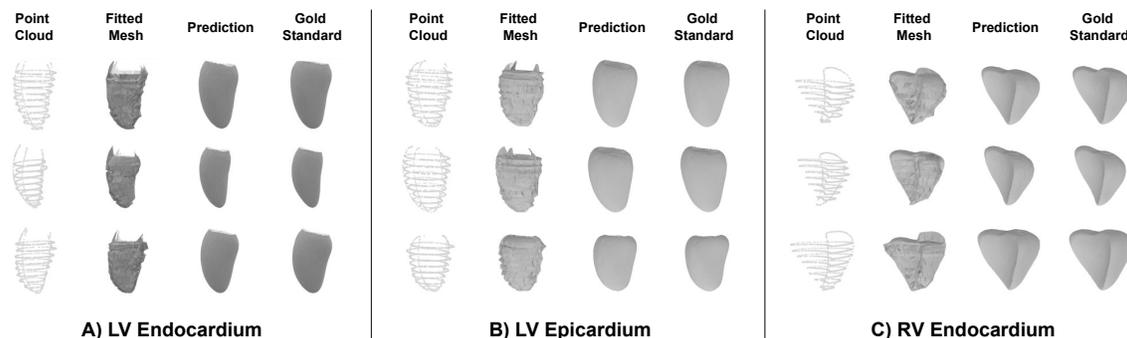


Figure 3: Reconstruction results of three sample cases from the SSM dataset.

We observe that the mesh fitting step has successfully incorporated the anatomical information of the sparse point clouds while maintaining a realistic shape despite the shared

connectivity restriction. The predicted surface meshes align well with the respective gold standard ones on both a global and local level and for different input shapes and misalignments. The performance is consistent across all three cardiac substructures.

After a qualitative validation, we next aim to quantify the reconstruction performance of the Mesh Deformation U-Net on the SSM dataset. We select the median surface distance (SD), the mean SD, and the Hausdorff distance between the respective predicted and gold standard meshes of the test dataset as our evaluation metrics and report the results separately for each cardiac substructure in Table 1.

Table 1: Reconstruction results of the Mesh Deformation U-Net on the SSM dataset.

Cardiac structure	Median SD (mm)	Mean SD (mm)	Hausdorff (mm)
LV endocardium	0.78 (± 0.17)	0.98 (± 0.23)	4.69 (± 1.67)
LV epicardium	0.77 (± 0.19)	1.00 (± 0.25)	4.65 (± 1.41)
RV endocardium	0.97 (± 0.25)	1.25 (± 0.32)	4.77 (± 1.46)

Values represent mean/median (\pm standard deviation/quartile deviation).

We find surface distance values below the voxel size of the underlying 3D MRI acquisition ($1.25 \times 1.25 \times 2$ mm) for all cardiac substructures, with slightly lower scores for the left ventricular reconstruction than the right ventricular one.

Next, we compare the performance of the Mesh Deformation U-Net with various benchmark methods on the SSM dataset (Table 2). As our baseline approach without any deep learning component, we select Laplacian Mesh Smoothing (Vollmer et al., 1999). Furthermore, we choose a multi-layer graph convolutional network (GCN) based on spatial graph convolutions (Kipf and Welling, 2016) and a multi-layer network with spectral graph convolutions (Defferrard et al., 2016) but without U-Net-inspired pooling and long skip connections as additional comparative benchmarks.

Table 2: Comparison of Mesh Deformation U-Net performance with benchmark methods.

Method	Median SD (mm)	Mean SD (mm)	Hausdorff (mm)
Laplacian Smoothing	1.32 (± 0.40)	1.77 (± 0.47)	12.00 (± 3.80)
Spatial GCN	0.92 (± 0.34)	1.24 (± 0.39)	8.26 (± 3.11)
Spectral GCN	0.83 (± 0.20)	1.05 (± 0.24)	6.51 (± 2.15)
Mesh Deformation U-Net	0.78 (± 0.17)	0.98 (± 0.23)	4.69 (± 1.67)

Values represent mean/median (\pm standard deviation/quartile deviation).

We find that the Mesh Deformation U-Net outperforms both the baseline and geometric deep learning benchmarks across all evaluation metrics.

3.2 Reconstruction Pipeline for Real MR Images

After the evaluation of the Mesh Deformation U-Net on synthetic MRI contours, we next assess the capability of our proposed 4-step pipeline to reconstruct 3D shapes from real MR

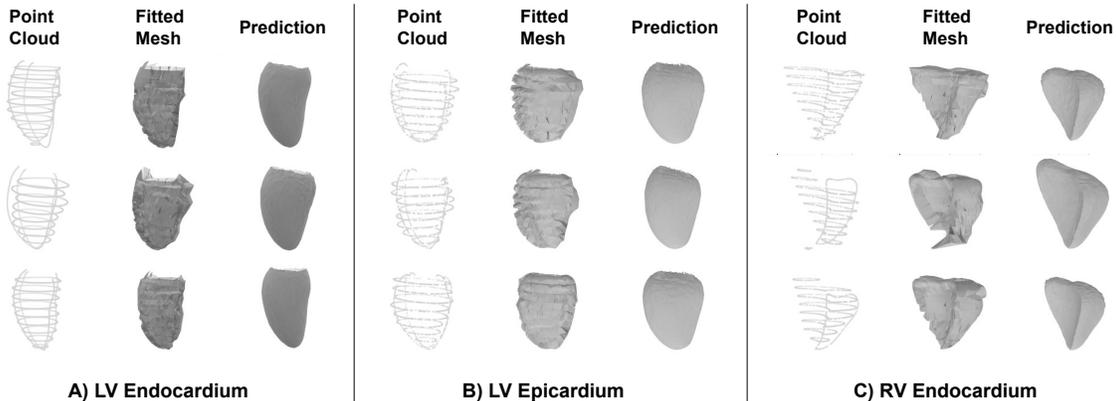


Figure 4: Reconstruction results of three sample cases of the UK Biobank dataset.

images. To this end, we select the Mesh Deformation U-Net pre-trained on the SSM dataset and apply it together with the three other pipeline steps to the raw MRI acquisition of the UK Biobank dataset. We execute the pipeline separately for each cardiac substructure and visualise the point cloud contours, approximately fitted mesh, and the final predicted mesh of three sample cases in Fig. 4.

Similar to the results on the SSM dataset, we observe that the mesh fitting step is sufficiently well designed to accurately encapsulate information from the anatomy contours in the fitted mesh, even for cases with stronger misalignment. The final mesh predictions exhibit realistic cardiac surfaces that adequately reflect the underlying input anatomies without showing any major remaining misalignment artifacts.

Since no gold standard 3D meshes are available for the UK Biobank dataset, we evaluate our pipeline’s reconstructions with multiple widely used clinical metrics on a population level. To this end, we calculate the LV volume, LV mass, and RV volume of the reconstructed UK Biobank meshes separately for female and male subjects and compare the results with two other cardiac MRI datasets derived from large-scale population studies in Table 3. While Petersen et al. (2017) computed the metrics based on 2D MRI slices using the modified Simpson’s rule, Bai et al. (2015) obtained the values directly from the 3D MRI acquisitions.

Table 3: Comparison of clinical metrics with two large-scale population studies.

Sex	Clinical Metric	Petersen et al. (2017)	Bai et al. (2015)	Mesh Deformation U-Net
Female	LV Volume (ml)	124 (± 21)	138 (± 24)	129 (± 21)
	LV Mass (g)	70 (± 13)	96 (± 16)	84 (± 18)
	RV Volume (ml)	130 (± 24)	-	131 (± 30)
Male	LV Volume (ml)	166 (± 32)	178 (± 36)	166 (± 38)
	LV Mass (g)	103 (± 21)	128 (± 24)	122 (± 29)
	RV Volume (ml)	182 (± 36)	-	174 (± 43)

Values represent mean (\pm standard deviation) in all cases.

We find that the values achieved by our pipeline generally lie between the ones reported by Petersen et al. (2017) and Bai et al. (2015), for all metrics and both subpopulations.

4. Discussion and Conclusion

In this work, we have developed the Mesh Deformation U-Net as the key component of a novel 3D cardiac surface reconstruction pipeline from cine MR images. Our experiments on the SSM dataset have shown that the Mesh Deformation U-Net is capable of accurately reconstructing hearts with errors below the underlying image resolution and outperforms various geometric deep learning benchmarks. On the one hand, this indicates that the proposed mesh fitting step is able to successfully retain the important anatomy information of the inputs, while simultaneously maintaining the same vertex connectivity and overcoming the sparsity issue in the point cloud contours. On the other hand, it demonstrates that the hierarchical architecture of the Mesh Deformation U-Net is adequately designed to directly process high resolution 3D mesh data in an efficient manner and capture relevant shape features on both a global and local level. These conclusions are further corroborated by the findings that the high reconstruction accuracy is achieved for three separate cardiac substructures and for a variety of different cardiac shapes and sizes while using highly sparse contours with varying misalignments as inputs. Furthermore, its outperformance of other geometric deep learning approaches with flat architectures indicates that the U-Net-inspired multi-scale structure is beneficial for cardiac surface reconstruction. In addition, we observe that the proposed pipeline reconstructs realistic 3D hearts on the real UK Biobank dataset with clinical metrics in line with other large-scale population studies. This shows not only that the SSM dataset is adequately designed to reflect real acquisition conditions, but also that both the mesh fitting step and the Mesh Deformation U-Net are highly suitable for cardiac surface reconstruction from real MR images. Their straightforward integration into the multi-step end-to-end pipeline also demonstrates a high degree of applicability with fast and fully automatic execution after training as well as high robustness, especially in light of the cross-domain transfer from synthetic to real data.

The proposed pipeline is also designed to simultaneously consider anatomy information from both SAX and LAX views, which is in contrast to many previous works focusing only on SAX slices (Lamata et al., 2014). This combined multi-view processing allows the network to utilise additional information and is especially beneficial in our highly sparse settings. It is also of particular importance in the apical and basal regions of the heart, where only very limited information is available from the SAX slices alone. However, due to the shape prior induced by the template mesh in our mesh fitting step and the direct mesh processing of our network, we hypothesise that our pipeline would still yield positive results even in the absence of LAX information or the presence of other slice information.

Due to the deep learning basis of our method, it comes with considerable speed and memory advantages after training, compared to case-specific reconstruction techniques (Banerjee et al., 2021a,b; Lamata et al., 2014), while retaining high levels of accuracy. These characteristics are crucial for an easy and robust scaling of the approach to large population cohorts. The benefits also still hold, although to a lesser extent, in comparison with previous voxelgrid-based deep learning approaches (Xu et al., 2019). Since the Mesh Deformation U-Net utilises geometric deep learning operations which can be directly applied to lightweight

mesh representations, it can process the same 3D cardiac surface data with a fraction of the time and memory requirements of a voxelgrid-based approach. This allows an easier scaling to both larger numbers of cases and higher data resolutions. It also enables the training of our network on a standard CPU as opposed to the typically required high-powered GPUs. While previous point cloud-based deep learning approaches (Beetz et al., 2021a) also share these advantages, they are not able to easily create meshes with correct topology and shared vertex-connectivity that are a prerequisite for many follow-up 3D cardiac modeling tasks. Furthermore, while the mesh fitting step of our pipeline induces additional complexity and possible errors, we believe that its regularising role as a geometric shape prior as well as its non-linear deformation abilities are still advantageous, especially in our highly sparse settings.

Acknowledgments

This research has been conducted using the UK Biobank Resource under Application Number ‘40161’. The authors express no conflict of interest. The work of M. Beetz is supported by the Stiftung der Deutschen Wirtschaft (Foundation of German Business). A. Banerjee is a Royal Society University Research Fellow and is supported by the Royal Society Grant No. URF\R1\221314. The work of A. Banerjee and V. Grau is supported by the British Heart Foundation (BHF) Project under Grant HSR01230. The work of V. Grau is also supported by the CompBioMed 2 Centre of Excellence in Computational Biomedicine (European Commission Horizon 2020 research and innovation programme, grant agreement No. 823712).

References

- Wenjia Bai, Wenzhe Shi, Antonio de Marvao, Timothy JW Dawes, Declan P O’Regan, Stuart A Cook, and Daniel Rueckert. A bi-ventricular cardiac atlas built from 1000+ high resolution MR images of healthy subjects and an analysis of shape and motion. *Medical Image Analysis*, 26(1):133–145, 2015.
- Wenjia Bai, Matthew Sinclair, Giacomo Tarroni, Ozan Oktay, Martin Rajchl, Ghislain Vialant, Aaron M Lee, Nay Aung, Elena Lukaschuk, Mihir M Sanghvi, et al. Automated cardiovascular magnetic resonance image analysis with fully convolutional networks. *Journal of Cardiovascular Magnetic Resonance*, 20(65):1–12, 2018.
- Abhirup Banerjee, Julià Camps, Ernesto Zacur, Christopher M. Andrews, Yoram Rudy, Robin P. Choudhury, Blanca Rodriguez, and Vicente Grau. A completely automated pipeline for 3D reconstruction of human heart from 2D cine magnetic resonance slices. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 379(2212):20200257, 2021a.
- Abhirup Banerjee, Ernesto Zacur, Robin P. Choudhury, and Vicente Grau. Optimised misalignment correction from cine MR slices using statistical shape model. In *Medical Image Understanding and Analysis*, pages 201–209. Springer International Publishing, 2021b.

- Abhirup Banerjee, Ernesto Zacur, Robin P. Choudhury, and Vicente Grau. Automated 3D whole-heart mesh reconstruction from 2D cine MR slices using statistical shape model. In *44th Annual International Conference of the IEEE Engineering in Medicine Biology Society (EMBC)*, pages 1702–1706, 2022.
- Marcel Beetz, Abhirup Banerjee, and Vicente Grau. Biventricular surface reconstruction from cine MRI contours using point completion networks. In *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, pages 105–109. IEEE, 2021a.
- Marcel Beetz, Abhirup Banerjee, and Vicente Grau. Generating subpopulation-specific biventricular anatomy models using conditional point cloud variational autoencoders. In *International Workshop on Statistical Atlases and Computational Models of the Heart*, pages 75–83. Springer, 2021b.
- Marcel Beetz, Julius Ossenbeng-Engels, Abhirup Banerjee, and Vicente Grau. Predicting 3D cardiac deformations with point cloud autoencoders. In *International Workshop on Statistical Atlases and Computational Models of the Heart*, pages 219–228. Springer, 2021c.
- Marcel Beetz, Abhirup Banerjee, and Vicente Grau. Multi-domain variational autoencoders for combined modeling of MRI-based biventricular anatomy and ECG-based cardiac electrophysiology. *Frontiers in Physiology*, page 991, 2022a.
- Marcel Beetz, Abhirup Banerjee, Yuling Sang, and Vicente Grau. Combined generation of electrocardiogram and cardiac anatomy models using multi-modal variational autoencoders. In *IEEE 19th International Symposium on Biomedical Imaging (ISBI)*, pages 1–4, 2022b.
- Xiang Chen, Nishant Ravikumar, Yan Xia, Rahman Attar, Andres Diaz-Pinto, Stefan K Piechnik, Stefan Neubauer, Steffen E Petersen, and Alejandro F Frangi. Shape registration with learned deformations for 3D shape reconstruction from sparse and incomplete point clouds. *Medical Image Analysis*, page 102228, 2021.
- Jorge Corral Acero, Andreas Schuster, Ernesto Zacur, Torben Lange, Thomas Stiermaier, Sören J Backhaus, Holger Thiele, Alfonso Bueno-Orovio, Pablo Lamata, Ingo Eitel, et al. Understanding and improving risk assessment after myocardial infarction using automated left ventricular shape analysis. *JACC: Cardiovascular Imaging*, 2022.
- Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, page 3844–3852, 2016.
- Maxime Di Folco, Pamela Mocerì, Patrick Clarysse, and Nicolas Duchateau. Characterizing interactions between cardiac shape and deformation by non-linear manifold learning. *Medical Image Analysis*, 75:102278, 2022.
- Matthias Fey and Jan E. Lenssen. Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.

- Thomas Joyce, Stefano Buoso, Christian T. Stoeck, and Sebastian Kozerke. Rapid inference of personalised left-ventricular meshes by deformation-based differentiable mesh voxelization. *Medical Image Analysis*, 79:102445, 2022.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- Pablo Lamata, Matthew Sinclair, Eric Kerfoot, Angela Lee, Andrew Crozier, Bojan Blazevic, Sander Land, Adam J Lewandowski, David Barber, Steve Niederer, et al. An automatic service for the personalization of ventricular cardiac meshes. *Journal of The Royal Society Interface*, 11(91):20131023, 2014.
- Lei Li, Julia Camps, Abhirup Banerjee, Marcel Beetz, Blanca Rodriguez, and Vicente Grau. Deep computational model for the inference of ventricular activation properties. *arXiv preprint arXiv:2208.04028*, 2022.
- Charlène Mauger, Kathleen Gilbert, Avan Suinesiaputra, Beau Pontre, Jeffrey H Omens, Andrew D McCulloch, and Alistair A Young. An iterative diffeomorphic algorithm for registration of subdivision surfaces: application to congenital heart disease. In *40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 596–599. IEEE, 2018.
- Charlène Mauger, Kathleen Gilbert, Aaron M Lee, Mihir M Sanghvi, Nay Aung, Kenneth Fung, Valentina Carapella, Stefan K Piechnik, Stefan Neubauer, Steffen E Petersen, et al. Right ventricular shape and function: cardiovascular magnetic resonance reference morphology and biventricular risk factor morphometrics in UK Biobank. *Journal of Cardiovascular Magnetic Resonance*, 21(1):1–13, 2019.
- Kate McLeish, Derek L. G. Hill, David Atkinson, Jane M. Blackall, and Reza Razavi. A study of the motion and deformation of the heart due to respiration. *IEEE Transactions on Medical Imaging*, 21(9):1142–1150, 2002.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. PyTorch: An imperative style, high-performance deep learning library. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, page 8026–8037, 2019.
- Steffen E Petersen, Paul M Matthews, Fabian Bamberg, David A Bluemke, Jane M Francis, Matthias G Friedrich, Paul Leeson, Eike Nagel, Sven Plein, Frank E Rademakers, et al. Imaging in population science: cardiovascular magnetic resonance in 100,000 participants of UK Biobank - rationale, challenges and approaches. *Journal of Cardiovascular Magnetic Resonance*, 15(46):1–10, 2013.
- Steffen E Petersen, Nay Aung, Mihir M Sanghvi, Filip Zemrak, Kenneth Fung, Jose Miguel Paiva, Jane M. Francis, Mohammed Y. Khanji, Elena Lukaschuk, Aaron M. Lee, et al. Reference ranges for cardiac structure and function using cardiovascular magnetic

- resonance (CMR) in Caucasians from the UK Biobank population cohort. *Journal of Cardiovascular Magnetic Resonance*, 19(18):1–19, 2017.
- Anurag Ranjan, Timo Bolkart, Soubhik Sanyal, and Michael J Black. Generating 3D faces using convolutional mesh autoencoders. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 704–720, 2018.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- Michael Benjamin Stokes and Ross Roberts-Thomson. The role of cardiac imaging in clinical practice. *Australian prescriber*, 40(4):151, 2017.
- Yi Su, May-Ling Tan, Chi-Wan Lim, Soo-Kng Teo, Senthil Kumar Selvaraj, Min Wan, Liang Zhong, and Ru-San Tan. Automatic correction of motion artifacts in 4D left ventricle model reconstructed from MRI. In *Computing in Cardiology*, pages 705–708, 2014.
- Avan Suinesiaputra, Pierre Ablin, Xenia Alba, Martino Alessandrini, Jack Allen, Wenjia Bai, Serkan Cimen, Peter Claes, Brett R Cowan, Jan D’hooge, et al. Statistical shape modeling of the left ventricle: myocardial infarct classification challenge. *IEEE Journal of Biomedical and Health Informatics*, 22(2):503–515, 2017.
- Benjamin Villard, Ernesto Zacur, Erica Dall’Armellina, and Vicente Grau. Correction of slice misalignment in multi-breath-hold cardiac MRI scans. In *Statistical Atlases and Computational Models of the Heart. Imaging and Modelling Challenges*, pages 30–38, 2017.
- Benjamin Villard, Vicente Grau, and Ernesto Zacur. Surface mesh reconstruction from cardiac MRI contours. *Journal of Imaging*, 4(1):16, 2018.
- Jörg Vollmer, Robert Mencl, and Heinrich Mueller. Improved laplacian smoothing of noisy surface meshes. In *Computer graphics forum*, volume 18, pages 131–138. Wiley Online Library, 1999.
- Zhao-Yang Wang, Xiao-Yun Zhou, Peichao Li, Celia Theodoreli-Riga, and Guang-Zhong Yang. Instantiation-net: 3D mesh reconstruction from single 2D image for right ventricle. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 680–691. Springer, 2020.
- Hao Xu, Ernesto Zacur, Jurgen E Schneider, and Vicente Grau. Ventricle surface reconstruction from cardiac MR slices using deep learning. In *International Conference on Functional Imaging and Modeling of the Heart*, pages 342–351. Springer, 2019.
- Xiao-Yun Zhou, Zhao-Yang Wang, Peichao Li, Jian-Qing Zheng, and Guang-Zhong Yang. One-stage shape instantiation from a single 2D image to 3D point cloud. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 30–38. Springer, 2019.