SB^{3}C2023
Summer Biomechanics, Bioengineering and Biotransport Conference
June 4–8, 2023, Vail, CO, USA

LEARNING DIFFEOMORPHIC DEFORMATIONS FOR WHOLE HEART MESH GENERATION

Arjun Narayanan[1], Fanwei Kong[2], Shawn C. Shadden[1]

[1]Department of Mechanical Engineering, University of California Berkeley, Berkeley, California, USA.
[2]Department of Pediatrics, Stanford University, Stanford, California, USA.

INTRODUCTION

Image-based computer modeling is playing an increasing role in understanding the mechanisms of cardiac disease and personalized care.[1] Broadly, this paradigm uses medical imaging (CT or MR) to construct an anatomically accurate computer model of the heart in order to mathematically model physiological processes and probe functional information. Reconstructing an accurate, personalized computer model of the heart is challenging because of imaging artifacts, limited resolution and difficultly differentiating between cardiac and surrounding tissues.

Segmentation is the process of identifying structures of interest in an image. Machine learning has been utilized to automate segmentation of cardiac structures from CT or MR images.[2] However, segmentation often generates artifacts that are unfit for simulation-based modeling. Recently, we have developed alternative template based deep-learning methods that are able to generate simulation-ready computer models of cardiac structures automatically from cardiac images.[3] However, these are mostly surface models of the blood pool boundaries since only these boundaries can be discerned from clinical imaging. Developing models wherein *tissue thickness* is assigned to all cardiac structures has remained challenging since most tissues are too thin to be accurately discerned.

We present here a deep learning method to produce whole-heart meshes with tissue thickness from clinical image data. We employ a template-based method to ensure accurate, simulation compatible models. Briefly, we use deep learning to deform a template mesh using a combination of linear-transformations and diffeomorphic flow deformations.[4] We present a novel physics-informed loss term that penalizes vanishing volumes thereby preventing self-penetrations in thin walled structures. The model can be trained on data containing no thickness information and subsequently evaluated on a template with thickness. We demonstrate that this approach is able to successfully deform this new template in a realistic manner without interpenetration. The predictions of our model can be readily used to generate full 3D tetrahedral meshes.

METHODS

Dataset information. We used the multi-modality whole heart segmentation challenge (MMWHS)[2] dataset. We split the dataset into 16 CT and 16 MR samples for training and used the remaining 4 CT and 4 MR samples for validation. We augmented the training dataset by introducing small perturbations including random scaling, translation, rotation, shear and local b-spline deformations to produce a final dataset size of 1920 samples.

Neural network architecture. We observed a large difference in scale and position across the data-set. This can be attributed to the different field-of-view in the image samples along with inter-patient variations. Hence, our framework consists of two modules (1) a 3D convolutional neural network (CNN) that performs a linear transformation on the template, and (2) a 3D Unet that learns to deform the template by integrating mesh vertices along a diffeomorphic flow vector field. The linear transformation is effective at scaling and positioning the template, and the flow deformation captures the finer details. For memory reasons the modules are trained separately. We use a weighted sum of the 2-way chamfer distance and geometrical regularizations such as normal consistency, mesh edge loss, and laplace loss as objective. We introduce a novel physics-informed loss term that penalizes vanishing volumes. A reference volume V_0 at time t_0 evolves to a volume V_1 at time t_1 in a flow vector field f as,

$$\frac{V_0}{V_1} = \exp\left(-\int_{t_0}^{t_1} \operatorname{div}(f)\ \mathrm{dt}\right) \tag{1}$$

If a finite initial volume V_0 collapses to an infinitesimal volume V_1 after deformation, equation 1 evaluates to a very large value. We incorporate this term into the objective function to discourage collapsing volumes. The divergence of the flow field is computed using a central finite difference scheme. Fig. 1 shows a schematic of our network architecture.

RESULTS

Dice scores. We use 4 CT and 4 MR samples from the MMWHS training data for validation. The predicted final meshes are converted to segmentations. We compute the dice score between predicted and ground truth segmentations which is presented in Table 1. The predictions of our model for a typical validation sample are shown in Fig. 2.

Self-intersections. We measured self-intersections using PyMesh. The predicted final meshes contained *zero* self-intersections within cardiac structures. For the predictions to

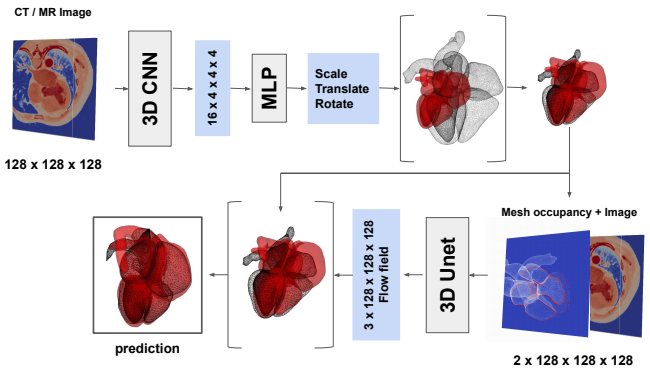

Figure 1: 128^3 input image processed by 3D CNN to produce 4×16^3 encoding. Linear network predicts 3 scale, 3 translate, and 3 rotate parameters from encoding to transform the template mesh. Transformed vertices embedded in an occupancy map and passed with image as input to 3D Unet which generates 3D flow vector field. Vertices are integrated along flow field to produce final prediction.

Table 1: Median and minimum dice scores on validation data for the converged model.

Cardiac structure	CT		MR	
	Median	Min	Median	Min
Myo	0.83	0.78	0.80	0.72
LA	0.91	0.87	0.83	0.75
LV	0.91	0.82	0.94	0.88
RA	0.89	0.86	0.84	0.76
RV	0.90	0.85	0.89	0.85
Ao	0.94	0.86	0.83	0.70
PA	0.73	0.63	0.70	0.10

be intersection-free it is important that the initial template is also intersection-free.

Patient-specific myocardium. An advantage of our framework is that once the network is trained, different templates can be used without re-training. For example, Fig. 3(a) shows a myocardium template that was produced by fusing the myocardium mesh with parts of the left atria and aorta and adding thickness to these openings. It is important that these regions have thickness because (a) they are used for solid modeling in FSI simulations (b) we apply boundary conditions on these surfaces. Note that this template contains about 47k triangular faces which is significantly more than the template used for training. Fig. 3(b) shows a patient-specific prediction for a model trained without the loss term eq. 1. This model collapses the thickness of the mitral and aortic tissue into an unphysical thin surface. It is not possible to generate a 3D volumetric mesh from such a surface mesh. The worst case predicted mesh contained 234 self-intersecting faces. Fig. 3(c) shows a patient-specific prediction for a model trained with the loss term eq. 1. This model preserves the thickness of the mitral and aortic openings. The worst case predicted mesh contained 8 self-intersecting faces which is a 97% improvement. The remaining self-intersections can be readily eliminated by running a constrained smoothing algorithm. Subsequently, we were able to generate full 3D tetrahedral meshes from these surface meshes.

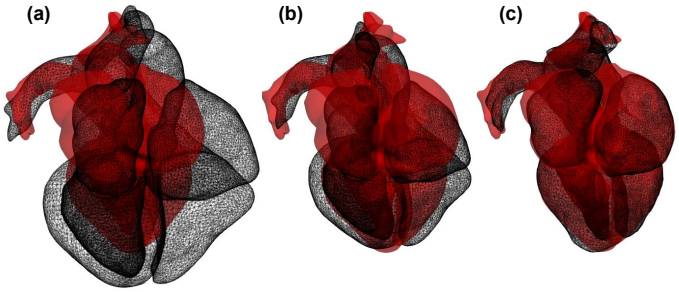

Figure 2: Predictions of the model for a typical validation sample. Red surface is the ground truth and black wireframe is the template. (a) Undeformed initial template (b) after linear transformation (c) after flow deformation. The final average dice score was 0.86 for this sample.

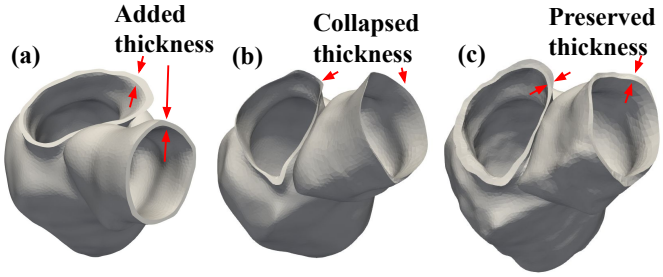

Figure 3: (a) Initial myocardium template with added thickness for mitral and aortic openings (b) prediction of model trained without divergence loss collapses thin walls resulting in self-intersections (c) prediction of model with divergence loss preserves thickness.

DISCUSSION

We have developed an automated deep learning method for generating patient-specific meshes of the human heart directly from biomedical images. These meshes are anatomically accurate and physically realistic. We presented a novel physics-informed loss term that penalizes collapsing volumes. This enables us to represent cardiac structures with tissue thickness even when this information is missing in the image data. Our predictions are almost entirely free of self-intersections, and the few that remain can be fixed by simple smoothing procedures. Our model predictions can be used to generate full 3D tetrahedral meshes with tissue thickness. We hope to soon present results of our model trained on a larger training dataset, provide detailed comparisons of our results against prior state of the art, and utilization of this method to applications of cardiac physiology simulation.

ACKNOWLEDGMENTS

This work was supported by NSF Award 1663671.

REFERENCES

[1] Prakosa et al. In: *Nature Biomed Eng.* 2.10 (2018), pp. 732–740. [2] Zhuang et al. In: *Med Image Analysis* 58 (2019), p. 101537. [3] Kong et al. In: *Med Image Analysis* 74 (2021), p. 102222. [4] Lebrat et al. In: *NeurIPS* 34 (2021), pp. 29491–29505.