

UC San Diego

UC San Diego Previously Published Works

Title

M-SiSSR: Regional Endocardial Function Using Multilabel Simultaneous Subdivision Surface Registration

Permalink

<https://escholarship.org/uc/item/9kr942z9>

ISBN

9783030787097

Authors

Vigneault, Davis M
Contijoch, Francisco
Bridge, Christopher P
et al.

Publication Date

2021

DOI

10.1007/978-3-030-78710-3_24

Peer reviewed



Published in final edited form as:

Funct Imaging Model Heart. 2021 June ; 12738: 242–252. doi:10.1007/978-3-030-78710-3_24.

M-SiSSR: Regional Endocardial Function Using Multilabel Simultaneous Subdivision Surface Registration

Davis M. Vigneault¹, Francisco Contijoch², Christopher P. Bridge³, Katherine Lowe², Chelsea Jan², Elliot R. McVeigh²

¹Department of Internal Medicine, Scripps Mercy Hospital, San Diego, CA, USA

²Department of Bioengineering, UC San Diego, La Jolla, CA, USA

³Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Charlestown, MA, USA

Abstract

Quantification of regional cardiac function is a central goal of cardiology. Multiple methods, such as Coherent Point Drift (CPD) and Simultaneous Subdivision Surface Registration (SiSSR), have been used to register meshes to the endocardial surface. However, these methods do not distinguish between cardiac chambers during registration, and consequently the mesh may “slip” across the interface between two structures during contraction, resulting in inaccurate regional functional measurements. Here, we present Multilabel-SiSSR (M-SiSSR), a novel method for registering a “labeled” cardiac mesh (with each triangle assigned to a cardiac structure). We compare our results to the original, label-agnostic version of SiSSR and find both a visual and quantitative improvement in tracking of the mitral valve plane.

Keywords

Cardiac Computed Tomography (CCT); Regional cardiac function; Convolutional Neural Networks (CNNs); Subdivision surfaces; Mesh registration

1 Introduction

Fully-automated quantification of regional cardiac function is a central goal of cardiology. Multiple methods, such as Coherent Point Drift (CPD) [13] and Simultaneous Subdivision Surface Registration (SiSSR) [19], have been used to register meshes to the endocardial surface of the left ventricle (LV). However, there is increasing interest in quantifying function beyond the LV, including the left atrium (LA). Although these methods could in principle be extended, naïve application to a multi-chamber mesh does not actually incorporate the additional information which the labels provide into the registration. Consequently, the mesh may “slip” across the interface between two structures during contraction, resulting in inaccurate regional functional measurements. Here, we present Multilabel-SiSSR (M-SiSSR), a novel extension of the SiSSR method which may be used to register a “labeled” cardiac mesh (where each cell is labeled according to its corresponding structure). We compare our results to the original, label-agnostic version of SiSSR, and find both a visual and quantitative improvement in tracking of the mitral valve plane.

2 Methods

2.1 Dataset and Annotation

100 3D+Time CCT scans were collected retrospectively from three institutions as part of an IRB-approved study. Segmentation of the first acquired and end systolic (ES) frames was performed using ITK-Snap [20,21]. The segmentation was initialized using spheres with radius 5 throughout the anatomy of interest, and the segmentation was allowed to proceed until the structure was filled. Region competition force and smoothing (curvature) force were left at the default values of 1.0 and 0.2, respectively. The left and right inferior and superior pulmonary veins (PVs), LA, left atrial appendage (LAA), LV, left ventricular outflow tract (LVOT), and ascending aorta (AA) were separately labeled. Planes separating adjacent structures were chosen using the ITK-Snap scalpel tool.

2.2 Left Heart Segmentation

The V-Net [10], an extension of the U-Net [14] for 3D segmentation, was chosen as the basic network for prediction of a segmentation \mathcal{S}_f from an input volume \mathcal{I}_f . This architecture consists of a convolutional layer, a batch-normalization layer, and a rectified linear unit (ReLU) repeated in sequence. Every second cycle, a max-pooling step or an interpolation step is inserted, along the downsampling- and upsampling-paths, respectively (four downsampling/upsampling steps total in our implementation). With each downsampling and upsampling step, the length of the feature vector was doubled and halved, respectively.

Data augmentation consisted of random uniform translation (up to 10% of the image width), rotation (up to 15°), and scaling (up to 10%). All networks were trained using five-fold cross validation for 96 epochs, and all images from a single patient were partitioned into the same cross-validation fold. The learning rate was initialized to 10^{-3} and decreased by a factor of 10 every 24 epochs. The loss function was the categorical cross entropy between the output of the final softmax layer and the ground truth segmentation. The batch size, initial feature map depth, and spacing to which the input images were isotropically resampled were treated as hyperparameters.

2.3 Boundary Candidate Selection

The first partition (20 subjects, 40 image volumes) was used for all subsequent experiments. Given a segmentation \mathcal{S}_f , we calculated the segmentation surface using the Cuberille algorithm [5]. In order to appropriately limit the scope of this work, only the immediately relevant labels (LA, LV, and LVOT) were retained. The boundary candidates \mathcal{E}_f were calculated as the centroids of the surface mesh cells.

2.4 Mesh Model Generation

Next we derived a template mesh \mathcal{T} from the mid systolic segmentation \mathcal{S}_{MS} ; the mid systolic phase was chosen so as to minimize the required displacement of the model during registration. The Cuberille algorithm was again used to derive an initial triangle mesh, and each cell was associated with the label of the adjacent foreground structure. Our application

requires that the extracted surface be a water-tight 2D manifold, and so the Cuberille algorithm was preferred over the ubiquitous Marching Cubes algorithm [9] due to the latter's well-known topological inconsistencies [4,11,12]. The resulting mesh was then decimated to a target number of cells using successive edge contraction. Two strategies were compared: midpoint decimation (in which the contracted edge is placed at the geometric mean of the two vertices) and Lindstrom-Turk decimation (in which the new vertex is placed so as to preserve the volume of the original model, [7]).

We found that naïve application of either method to a multi-label mesh resulted in distorted boundaries between structures and poor approximation of the initial mesh. To mitigate this effect, edge contraction was handled differently depending on the relationship between the candidate edge and the label boundaries, according to the following rules:

- “Interior” edges (both vertices surrounded by cells of the same label) and “boundary” edges (flanked by cells of differing labels) were contracted as usual, using the respective decimation technique.
- “Hanging” edges (one vertex lying on a structure boundary, the other surrounded by cells entirely of the same label) were contracted to the boundary vertex.
- “Bridging” edges (flanked by cells of the same label, but both vertices lying on structure boundaries) were not candidates for contraction.

2.5 Subdivision Surface Evaluation

Following the notation given in [19], let $\mathbf{X} \in \mathbb{R}^{3 \times (C \times F)}$ be the matrix of control point coordinates defining the sequence of template meshes \mathcal{T}_f , where C is the number of control points and F is the number of cardiac phases. Displacements from the elements of \mathbf{X} serve as the $3 \times C \times F$ parameters of the optimization. This control point matrix defines F triangular Loop subdivision surfaces, which we use to model the endocardial and endovascular surface. The position of a point $\mathbf{u} \in \mathbb{R}^3$ on the surface of \mathcal{T}_f may be calculated in terms of the matrix \mathbf{X} , the frame ℓ , the patch index i , and a set of parametric coordinates $\mathbf{t} = (t, s)$, which are local to a given patch. We define such a function $\chi: \mathbf{t}, \mathbf{X}_{f,i} \rightarrow \mathbb{R}^3$ as follows:

$$\chi(\mathbf{t}, \mathbf{X}_{f,i}) = \mathbf{X}_{f,i} \left(\mathbf{P}_k \bar{\mathbf{A}} \mathbf{A}^{n-1} \right)^\top \mathbf{b}(\hat{\mathbf{t}}). \quad (1)$$

Here n is the number of required subdivisions, \mathbf{A} is the *extended subdivision matrix*, $\bar{\mathbf{A}}$ is the *bigger subdivision matrix*, \mathbf{P}_k is the *picking matrix*, \mathbf{b} is the column of basis functions, and $\hat{\mathbf{t}}$ is a tuple representing the transformed patch coordinates. We refer the reader to [16,19] for precise definitions of these matrices.

2.6 Multilabel SiSSR (M-SiSSR)

Registration of the sequence of template meshes \mathcal{T} to the boundary candidates \mathcal{B} was performed using Levenberg-Marquardt least squares optimization. In the SiSSR algorithm, the primary residuals E_{cf} were the Cartesian components of the distance between the points

sampled from the surface of \mathcal{T}_f and the nearest point in the corresponding boundary candidate mesh \mathcal{C}_f :

$$E_{cf} = \sum_{f,i,t} \|\mathbf{u}_{f,i,t} - \phi(\mathbf{u}_{f,i,t})\|^2. \quad (2)$$

Here, $\phi: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ is the function which computes the nearest point in the boundary candidate mesh, and is represented internally as a vector of K_d trees. In M-SiSSR, we propose a modified cost function $E_{c\hat{f}}$ (relying on a modified function $\hat{\phi}: \mathbb{R}^3 \rightarrow \mathbb{R}^3$), which represents the nearest point *of the same label* in the corresponding boundary candidate mesh. Internally, this is represented as a vector of vectors of K_d trees (one for each label and cardiac phase).

Some amount of regularization is necessary to encourage biologically plausible surfaces and control point trajectories in time. In this work, we consider thin plate energy $E_{\hat{t}p}$ (which encourages surface smoothness in the second derivative), acceleration E_{ac} (which encourages smooth control point trajectories in time), and edge length E_{el} (which provides tension and discourages “foldover” artifacts), as defined in [19].

The overall loss is formulated as follows:

$$E = \min(\alpha_{c\hat{f}} E_{c\hat{f}} + \alpha_{\hat{t}p} E_{\hat{t}p} + \alpha_{ac} E_{ac} + \alpha_{el} E_{el})$$

Here the coefficients $\alpha_{c\hat{f}}$, $\alpha_{\hat{t}p}$, α_{ac} , and α_{el} are scaling factors, which were treated as hyperparameters (defined as in [19]).

2.7 Implementation

The CNN segmentation architecture was implemented in Python 3.6.8 using the Keras 2.3.1 interface to Tensorflow 1.15.0 [1] in an Ubuntu 18.04 Docker container. Models were trained on Amazon Web Services p2.xlarge instances with a single NVIDIA K80 GPU with 12GB memory. Mesh visualization was performed using VTK [15]; mesh decimation was performed using CGAL [18]; mesh registration was implemented using the Ceres Solver [2]; and the remaining mesh processing was performed using ITK [6].

3 Results

3.1 CNN Segmentation

The CNN was trained for a range of spatial resolutions and feature map depths. The best-performing network for the LA, LV, and LVOT was that with images resampled to 2mm isotropic with 32 channels in the initial feature map and trained with a batch size of 4; the predictions of this network were used for all subsequent experiments (Table 1).

3.2 Template Mesh Generation

When generating the template mesh, we target a specific number of faces in the decimation algorithm. As the number of faces increases, the computational time of the mesh registration

as well as the ability to represent fine anatomical detail increases, while the regional influence of a particular control point decreases. As expected, Jaccard index increased as a function of the target number of faces (Fig. 1) for both midpoint and Lindstrom-Turk decimation techniques.

Moreover, in terms of Jaccard index, the Lindstrom-Turk method outperformed the midpoint method for any fixed target number of faces. Note, for example, that even at 64 faces, the LVOT is preserved in the Lindstrom-Turk model (Fig. 2b), whereas it is entirely lost in the midpoint model (Fig. 2h). Note also that for models generated via the Lindstrom-Turk method, there was noticeable loss of detail around certain features such as the cardiac apex below 256 faces, but conversely there was minimal improvement either in Jaccard index or in subjective anatomic detail above approximately 256 faces. For these reasons, the 256-face Lindstrom-Turk model was chosen as the starting point for registration. Representative models are shown in Fig. 2.

As discussed above, naïve application of either midpoint or Lindstrom-Turk decimation algorithm to a multi-label mesh without careful preservation of the edge resulted in unacceptable distortion of the structure boundary (Fig. 3c). However, modification of the decimation procedure to avoid large displacements in the boundary edges resulted in a marked visual improvement (Fig. 3a).

3.3 Determining Scaling Factors

The coefficients of the regularization terms were explored in a stepwise fashion (Fig. 4), using Jaccard index and triangle condition (measures of segmentation accuracy and mesh quality, respectively) as outcomes. The goal is to maximize Jaccard index while minimizing triangle condition. The coefficients chosen for the final registration algorithm were $E_{\hat{p}} = 10^{-1}$, $E_{ac} = 10^{-1}$, and $E_{el} = 10^{-0.5}$.

3.4 Comparison to Label-Agnostic Approach: M-SiSSR vs SiSSR

As a comparison, the registration was run using the same regularization weight, but without accounting for the labels (i.e., using the label-agnostic SiSSR algorithm). Representative results are shown in Fig. 5. In the meshes registered through using the M-SiSSR method (Fig. 5a, 5d), the boundary between the LA and LV tracks that of the boundary candidates (Fig. 5b, 5e), excursing downward during systole. By contrast, in the meshes registered through using the SiSSR method (Fig. 5c, 5f), the boundary stays stationary, lying below and above that of the boundary candidates in the ED and ES frames, respectively.

Additionally, the centroid of the mitral valve annulus (approximated by the geometric mean of all points lying on the boundary between the LV and LA) was calculated at end diastole and end systole for the boundary candidate meshes, the meshes registered using the M-SiSSR method, and those registered using the SiSSR method. The error (as judged by the Euclidean distance between the boundary candidate and registered mesh centroids) was significantly lower for the M-SiSSR meshes compared to the SiSSR meshes both at end diastole (2.5 mm vs 3.9 mm, $p < 10^{-2}$) and at end systole (2.8 mm vs 5.6 mm, $p < 10^{-7}$).

We also calculated SQUEEZ (a metric of regional cardiac function used in CCT, [13]) from both the SiSSR and M-SiSSR methods. At left ventricular end systole, we expect there to be a sharp transition in SQUEEZ between the contracting LV and the filling LA. And in fact we do see such a sharp transition at the mitral valve plane in the M-SiSSR meshes (Fig. 6a), whereas in the SiSSR method, there is a gradual transition that appears to underestimate both LV contraction and LA expansion.

4 Discussion

In this work, we present a fully-automated pipeline in which a CNN is used to segment the left heart and connected vasculature from a sequence of CCT volumes; a labeled Loop subdivision surface is generated from the mid-systolic frame; and this model surface is registered to the CNN segmentations across all phases simultaneously, accounting for the label in the registration. The result is a patient-specific model which is topologically consistent across all timepoints, and therefore may be used to calculate regional function.

The SiSSR algorithm makes use of data structures and techniques described by Loop [8], Stam [16], Cashman [3], and Stebbing [17]. Compared to the first iteration of the SiSSR algorithm [19], the present work has a number of notable advantages and improvements. In particular, experiments were performed using human subjects rather than canine hearts, and the number of subjects included was higher (from 13 cases to 20); the segmentations were derived using a neural network, rather than through semi-automated level set segmentation; the LA and LVOT were included in the registration, rather than the LV alone; and the algorithm was modified to incorporate the identity of the underlying structure into the registration.

Although the number of subjects is increased from [19], the absolute number is still small; future work should be expanded to include the remaining 80 cases in the dataset. The LV myocardium, papillary muscles, and right heart are all of great interest, but are not included here; these should all be explored in future work. And although here we demonstrated a quantitative improvement in the tracking of the mitral valve compared to the label-agnostic SiSSR algorithm in terms of the error between the mitral valve centroids, in future work, it would be useful to additionally report an error metric with more immediate clinical relevance, such as mitral annular plane systolic excursion (MAPSE).

We hope that this work represents a substantive contribution to the use of subdivision surfaces for modeling of cardiac structures, and to the clinically important goal of quantifying regional cardiac function from CCT images.

Acknowledgments

This work is supported by National Institutes of Health grants R01HL144678 and K01HL143113.

References

1. Abadi M, et al.: TensorFlow: a system for large-scale machine learning (2016) <https://arxiv.org/pdf/1605.08695.pdf>
2. Agarwal S, Mierle K: Ceres Solver (2010) <http://ceres-solver.org>

3. Cashman TJ, Fitzgibbon AW: What shape are dolphins? building 3D morphable models from 2D images. *IEEE Trans. Pattern Anal. Mach. Intell* 35(1), 232–244 (2013). 10.1109/TPAMI.2012.68 [PubMed: 22392707]
4. Chernyaev E: Marching cubes 33: construction of topologically correct Isosurfaces. In: *GRAPHICON 1995*, Saint-Petersburg, Saint-Petersburg, Russia (1995)
5. Herman GT, Liu HK: Three-dimensional display of human organs from computed tomograms. *Comput. Graph. Image Process* 9(1), 1–21 (1979). 10.1016/0146-664X(79)90079-0. <https://linkinghub.elsevier.com/retrieve/pii/0146664X79900790>
6. Johnson HJ, Ibanez L, McCormick M: *The ITK Software Guide: Book 1*, 4 edn. (2016)
7. Lindstrom P, Turk G: Fast and memory efficient polygonal simplification. In: *Proceedings of the IEEE Visualization Conference*, pp. 279–286 (1998). 10.1109/visual.1998.745314
8. Loop C: Smooth subdivision surfaces based on triangles. Ph.D. thesis, University of Utah (1987) <http://www.citeulike.org/group/5490/article/2864922>
9. Lorensen WE, Cline HE: Marching cubes: a high resolution 3D surface construction algorithm. *ACM SIGGRAPH Comput. Graph* 21(4), 163–169 (1987). 10.1145/37402.37422. <http://portal.acm.org/citation.cfm?doid=37401.37422>
10. Milletari F, Navab N, Ahmadi SA: V-Net: fully convolutional neural networks for volumetric medical image segmentation. In: *CVPR* pp. 1–11 (2016). 10.1109/3DV.2016.79. <http://arxiv.org/abs/1606.04797>
11. Newman TS, Yi H: A survey of the marching cubes algorithm. *Comput. Graph* 30(5), 854–879 (2006). 10.1016/j.cag.2006.07.021
12. Nielson GM, Hamann B: The asymptotic decider: Resolving the ambiguity in marching cubes. In: *Proceedings of the 2nd Conference on Visualization 1991, VIS 1991* (1991). 10.1109/visual.1991.175782
13. Pourmorteza A, Schuleri KH, Herzka D.a., Lardo AC, McVeigh ER: A new method for cardiac computed tomography regional function assessment: Stretch quantifier for endocardial engraved zones (SQUEEZ). *Circulation Cardiovascular Imaging* 5(2), 243–250 (2012). 10.1161/CIRCIMAGING.111.970061 [PubMed: 22342945]
14. Ronneberger O, Fischer P, Brox T: U-Net: Convolutional Networks for Biomedical Image Segmentation. In: Navab N, Hornegger J, Wells WM, Frangi AF (eds.) *MICCAI 2015*. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). 10.1007/978-3-319-24574-4_28
15. Schroeder W, Martin K, Lorensen B: *The Visualization Toolkit*. Kitware, 4 edn. (2006)
16. Stam J: Evaluation of Loop subdivision surfaces. In: *SIGGRAPH*, pp. 1–15 (1998)
17. Stebbing R: *Model-Based Segmentation Methods for Analysis of 2D and 3D Ultrasound Images and Sequences*. DPhil, University of Oxford (2014)
18. The CGAL Project: *CGAL User and Reference Manual*. CGAL Editorial Board, 5.2 edn. (2020). <https://doc.cgal.org/5.2/Manual/packages.html>
19. Vigneault DM, Pourmorteza A, Thomas ML, Bluemke DA, Noble JA: SiSSR: simultaneous subdivision surface registration for the quantification of cardiac function from computed tomography in canines. *Med. Image Anal* 46, 215–228 (2018). 10.1016/j.media.2018.03.009 [PubMed: 29627686]
20. Yushkevich PA, et al. : User-guided segmentation of multi-modality medical imaging datasets with ITK-SNAP. *Neuroinformatics* 17(1), 83–102 (2018). 10.1007/s12021-018-9385-x
21. Yushkevich PA, et al. : User-guided 3D active contour segmentation of anatomical structures: significantly improved efficiency and reliability. *Neuroimage* 31, 1116–1128 (2006). 10.1016/j.neuroimage.2006.01.015 [PubMed: 16545965]

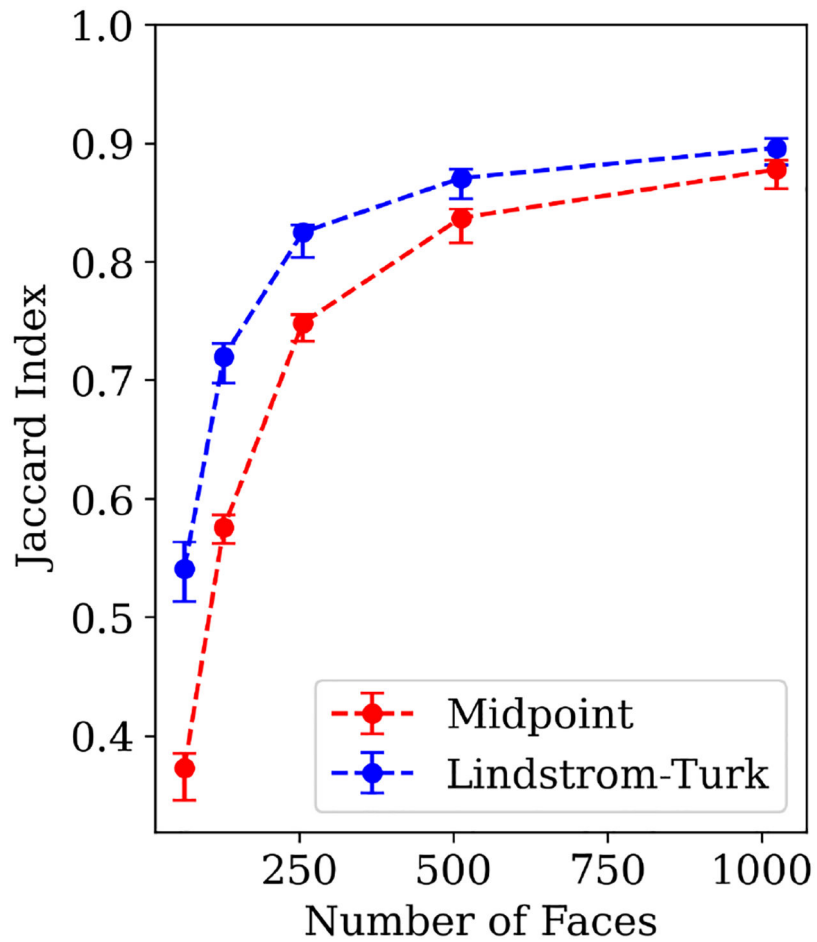


Fig. 1. Template mesh resolution vs Jaccard Index.

Jaccard index is plotted as a function of the number of faces targeted in the mesh decimation algorithm (either Midpoint or Lindstrom-Turk). Note that there is a precipitous decline in accuracy below 256 faces, and that the Lindstrom-Turk method outperforms the midpoint method at all resolutions.

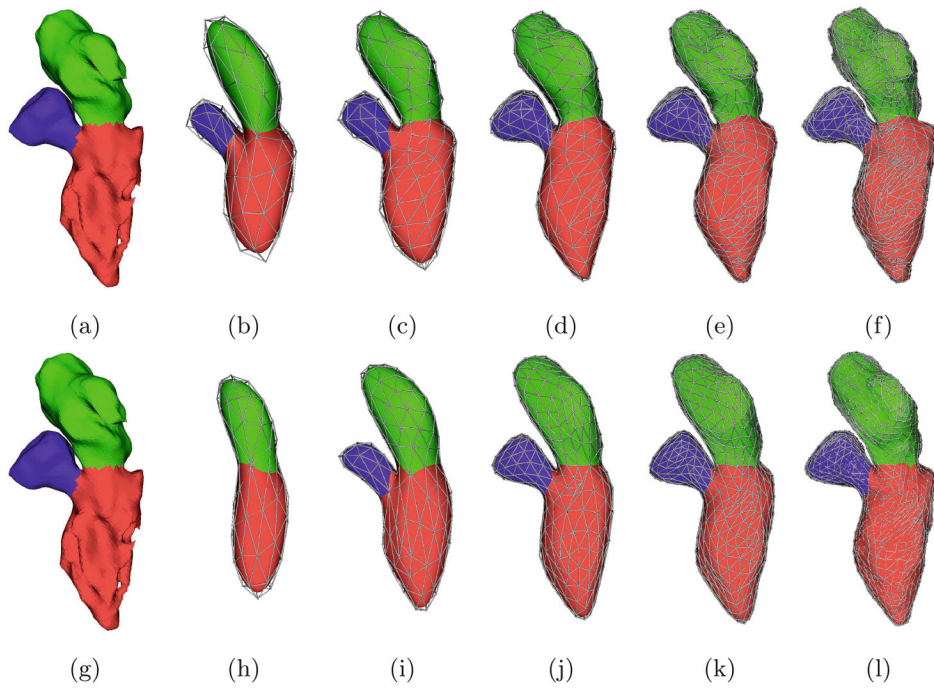


Fig. 2. Template mesh resolution.

Representative template meshes \mathcal{T} are shown at 64 (b, h), 128 (c, i), 256 (d, j), 512 (e, k), and 1024 (f, l) faces. The corresponding boundary candidate mesh \mathcal{E} (a, g) is also shown for comparison. In this and subsequent figures, the LA is green, the LV is red, and the LVOT is blue.




Fig. 3. Edge preservation.

Template meshes \mathcal{T} generated with (a) and without (c) edge preservation. The boundary candidates for the mid systolic frame are shown for comparison (b).

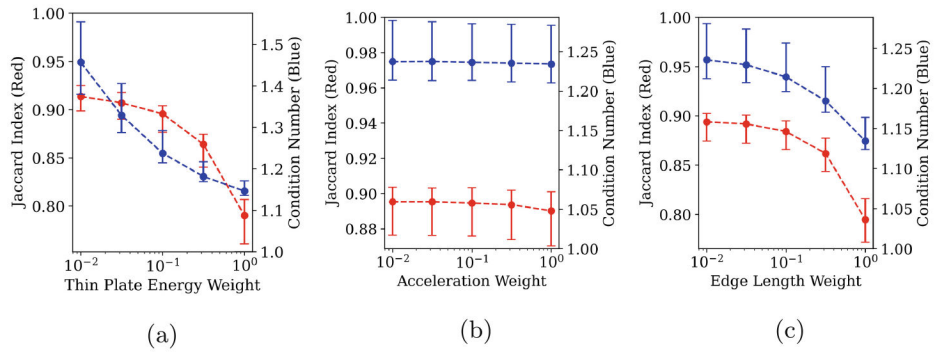


Fig. 4. Hyperparameter search.

(a) Thin plate energy parameter search, with $\alpha_{cf} = 1.0$. (b) Acceleration parameter search, with $\alpha_{cf} = 1.0$ and $\alpha_{ip} = 10^{-1}$. (c) Edge length parameter search, with $\alpha_{cf} = 1.0$, $\alpha_{ip} = 10^{-1}$, and $\alpha_{ac} = 10^{-1}$.

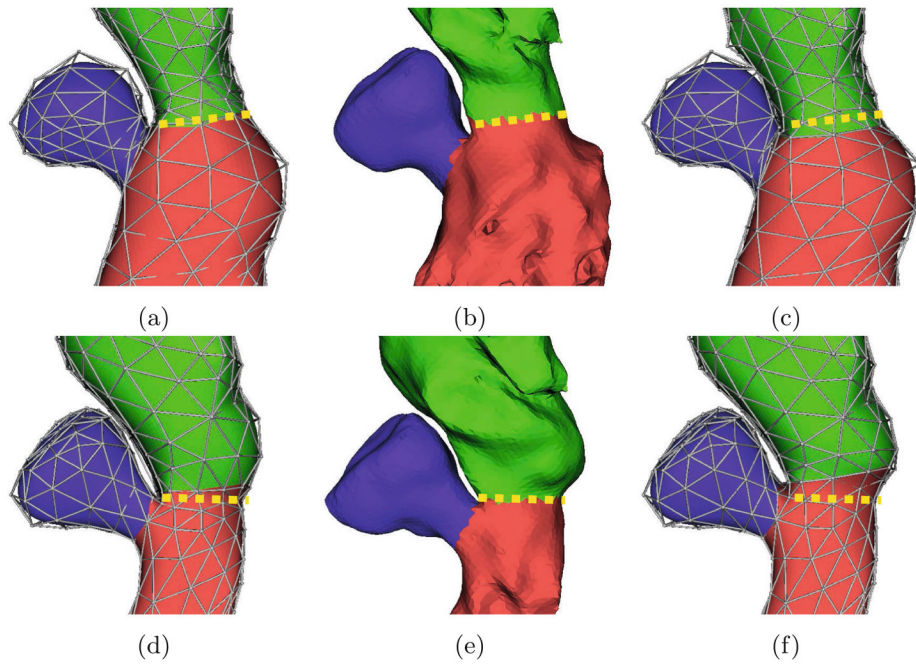


Fig. 5. Registered meshes.

Registered meshes using label information (a, d) and ignoring label information (c, f), and boundary candidates for comparison (b, e). Both ED (top) and ES (bottom) frames are shown. The mitral annular plane of the boundary candidate mesh is approximated by a dashed yellow line and reproduced over the M-SiSSR and SiSSR meshes. Note that the M-SiSSR meshes more closely track the mitral annular plane.




Fig. 6. SQUEEZ.

Representative meshes registered using (a) and ignoring (b) label information and colored according to their SQUEEZ values.

Table 1.
CNN segmentation performance.

Jaccard index for the LA, LV, and LVOT is reported as a function of image spacing, feature map depth, and batch size. The Jaccard index for the best performing network is bolded.

Spacing (mm)	Feature Map Depth	Batch Size	LV	LA	LVOT
1.0	8	1	71.9 ± 17.3	82.4 ± 12.0	63.6 ± 18.0
1.0	16	1	68.6 ± 21.7	74.7 ± 23.6	53.3 ± 26.7
2.0	8	14	75.3 ± 11.7	75.4 ± 9.8	0.4 ± 0.7
2.0	16	8	78.6 ± 10.0	84.1 ± 7.6	68.1 ± 8.8
2.0	32	4	79.2 ± 11.3	86.1 ± 6.5	70.8 ± 9.7
3.0	8	20	39.6 ± 24.0	42.2 ± 23.8	0.0 ± 0.0
3.0	16	12	68.2 ± 10.4	74.8 ± 6.3	42.9 ± 23.6
3.0	32	6	69.5 ± 9.6	76.1 ± 6.4	59.1 ± 9.4
3.0	64	3	69.6 ± 9.7	76.4 ± 6.4	59.7 ± 10.0

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript