

Cardiac Mesh Reconstruction from Sparse, Heterogeneous Contours

Benjamin Villard¹, Valentina Carapella¹, Rina Ariga², Vicente Grau¹, and Ernesto Zacur¹

¹ Institute of Biomedical Engineering, University of Oxford, United Kingdom,
`benjamin.villard@eng.ox.ac.uk`,

² University of Oxford Enter for Clinical Magnetic Resonance Research, Division of Cardiovascular Medicine, Radcliffe Department of Medicine, University of Oxford

Abstract. We introduce a tool to reconstruct a geometrical surface mesh from sparse, heterogeneous, non coincidental contours and show its application to cardiac data. In recent years much research has looked at creating personalised 3D anatomical models of the heart. These models usually incorporate a geometrical reconstruction of the anatomy in order to understand better cardiovascular functions as well as predict different processes after a clinical event. The ability to accurately reconstruct heart anatomy from MRI in three dimensions commonly comes with fundamental challenges, notably the trade off between data fitting and regularization. Most current techniques requires data to be either parallel, or coincident, and bias the final result due to prior shape models or smoothing terms. Our approach uses a smoothness term while maximizes the data fitting. Quantitative assessment of our method was performed on synthetic data obtained from a mean cardiac shape model, and qualitative assessment is presented on clinical data belonging to a normal and a hypertrophic cardiomyopathic subjects. Our method is both used on epicardial and endocardial left ventricle surfaces, but as well as on the right ventricle.

Keywords: Computational geometry, Mesh Reconstruction, Cardiac Modelling

1 Introduction

In recent years much research has looked at creating personalised 3D anatomical models of the heart [1–3]. These models usually incorporate a geometrical reconstruction of the anatomy in order to understand better cardiovascular functions as well as predict different processes after a clinical event. Besides being fundamental to shape analysis, they are the prerequisite to finite element analysis [4]. The ability to accurately reconstruct heart anatomy from MRI in three dimensions commonly comes with fundamental challenges, notably the trade off between data fitting and regularization. Most works in the literature tend to over regularize, or bias the data fitting process in order to obtain smooth meshes [5, 6].

Other methods work based on templates or shape prior which are usually learned on normal patient anatomies, and tend to fail when presented with pathological data [7]. Our approach differs in this, as our approach minimizes the use of a regularization term while maximizes the data fitting. In this paper we discuss a process to reconstruct a geometrical 3D surface mesh from a set of contours available from manual delineations as well as from automatic segmentations. The main contributions of this work are:

- Possibility to use our method on non parallel, sparse, heterogeneous contours
- Possibility to deal with non coincidental contours
- Possibility to use our technique on the right ventricle as well as the left ventricle

Cardiac contour are typically composed from a stack of short axes ranging from 8-10 slices, and a couple long axes. As such the input data is quite sparse for any conventional mesh reconstruction process such as isosurfacing [8]. In manual contouring, clinicians are seldom able to account for the 3D environment which can lead to discrepancies between long axes and short axes contours. This can occur, for example, due to papillary muscles being segmented in LAs, but not in SAs. Furthermore due to the image acquisition occurring at different breath hold, as well as any patient movement inside the scanner, the contours might be misaligned. This would further increase their discrepancies and extra steps might be needed to correct for this alignment [9]. This post-acquisition correction of the pose of the images can lead to losing the parallelism in the SA stack. Furthermore, due to the delineations not being able to grasp perfectly the overall shape of the heart, contour to contour distances cannot be minimized to obtain perfect coincidence and therefore spatial 3D consistency [9]. Finally, since many cardiomyopathies are localized in the left ventricle, many previous efforts were focused on methodologies specifically designed or tuned for the LV.

1.1 Current state of the art

Mesh construction is inherently problem specific. Whilst simple heart models can be used for mechanical studies, patient-specific meshes require more complex anatomical models [10]. These models are characterized though 3D meshes, on which shape analysis or finite element methods (FEM) can be performed. From a set of contours or delineated curves, a surface mesh can be rendered and act as an input to generate a volumetric mesh. Depending on the complexity of the mesh needed, many parameters need to be taken account for; notably if the visual appearance of the mesh needs to be smooth [2, 10]. In [2], the surface mesh is constructed by optimising over the topology, using a level-sets approach, and having the genus as a control parameter. Whilst this approach allows control over the topology, it can fail for sparse data. In order for the resulting mesh to represent the most accurate anatomical representation, an image volume is necessary. In order to obtain uniformly distributed contour points, the authors in [11, 12], first fit a cubic periodic B-spline curve to the

contours allowing them to uniformly sample candidate points. The cons of this approach is that it relies on the the number of control points selected which in turn smooths the overall mesh. This is similar to [6], which can cause for the resulting mesh to be highly convoluted when dealing with pathological data. The method in [13] is closest to our own, as it tries to maximize the data fitting, whilst smoothing the interpolated part of the mesh.

2 Materials and Methods

2.1 Data

Several datasets were used. Synthetic contours were created from an average MRI mesh obtained from [14, 15] and available at <http://wp.doc.ic.ac.uk/wbai/data/>. Using that mean shape, contours were synthesized by slicing the mesh at various spatial position. Parallel and non parallel contours were produced and small random translations and rotations applied to the contours. We treated the initial mean shape mesh as ground-truth and compared our results against it. Clinical contours were also used, which were delineated by experts on cine images at end-diastole. However as we do not have ground truth available for this data, we assess the resulting mesh by means of visual assessment only.

2.2 Initialization Mesh

Let $\{C_j\}$ be a set of contours, each lying on a plane with normal n_j . These normals are clustered and the mean normal of the cluster group containing the most frequent orientation of the planes is calculated and the transformation matrix to align it with the Z axis is determined. The resulting transformation is applied to all the contours such that the contours are in a common frame of reference aligned with respect to the Z axis. The contours with normals belonging to that cluster (which are expected to be the ones in SA view) are then sorted from the most apical to the most basal. Two extra points are added to $\{C_j\}$ to account for the upper and bottom lid of the initial mesh. We then proceed to create a rough initialization of the mesh by fitting a tubular surface going through the contour. For each of the contours, distribute K equally spaced contour points P_C . This is achieved by calculating the piecewise linear interpolation $P_C(K)$. A ruled surface approach is then used to obtain an initial mesh, where the initial triangle making up the mesh are composed of the set $\{(C_j(i), C_j(i+1), C_{j+1}(i+K)), (C_{j+1}(i), C_{j+1}(i+1), C_j(i+1))\}$. Depending on the availability of the LA information, the algorithm initializes the upper and lower bound mesh points as the mean position of the top and bottom most contour points lying on the plane passing through those points.

2.3 Mesh Deformation

Since the initial selected points make up the mesh, another set of contour points will be selected as the data fitting terms, and will act as attractors for the

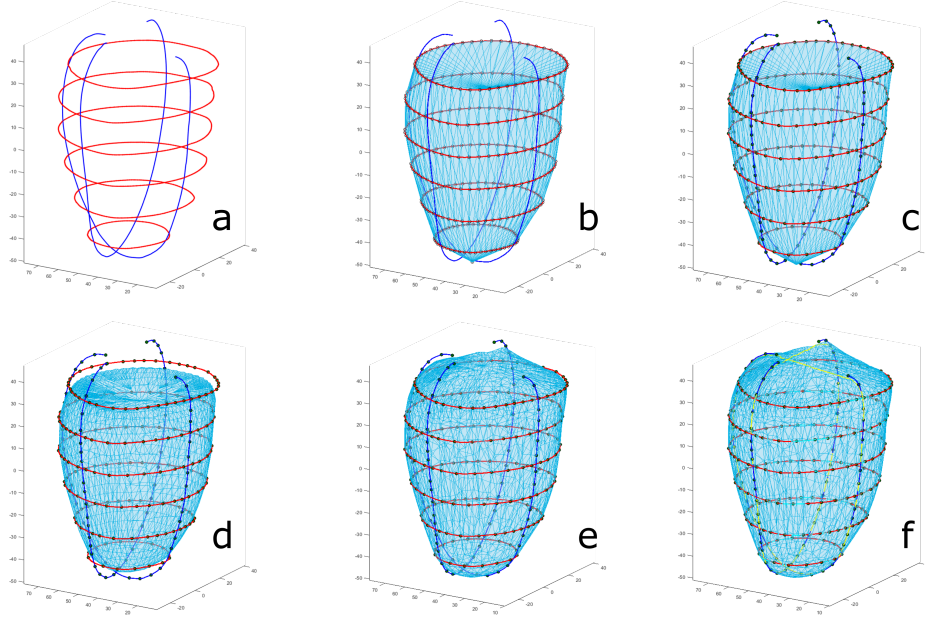


Fig. 1. Overall pipeline for generating a surface mesh from contours. a) The input contours, b) Initialization mesh, c-e) Different iterations of the process after deforming, subdividing and decimating the mesh, f) Final Mesh

mesh. N number of points are sampled, where $N \gg K$, using farthest point sampling, allowing for a representation of the global spatial distribution of the contour points. The furthest point method ensures that no contour is sampled more densely than any others, and enforces a global distribution of $\{C_j\}$. At this point, the initial mesh, the contours, and the attractors are all scaled to a common reference frame that is irrespective of the units of the provided contours. Once all the preprocessing is achieved, a deformation field is computed by taking the closest point from each of the attractor points to the mesh \mathcal{M} , and a dense force field is computed by using approximating thin plate splines [16]. The mesh is then deformed by using the inverse of the force field. Let $\{P_i\}$ be the set of closest points to $\{Q_i\}$ lying on the surface \mathcal{M} . Let be $F : \mathbb{R}^3 \Rightarrow \mathbb{R}^3$ such that:

$$\mathcal{M} = \min \lambda \sum_{i=1}^N (||F(P_i) - (Q_i)||)^2 + E_{TPS}(F) \quad (1)$$

$$E_{TPS}(f) = J_m^d(u)$$

where F is a mapping function and J the thin plate spline functional using derivatives of order m [17]. In this way our optimisation is a trade off between approximation and interpolation of the surface \mathcal{M} . Once the mesh has been deformed, the mesh undergoes subdivision and quadratic decimation. By subdividing the mesh we increase its resolution allowing for a more appropriate

reconstruction of the mesh. However as this process can lead to extremely high amounts of triangles, the mesh is therefore decimated, which allows some control over the amount of triangles while preserving its geometric characteristics [18]. Finally, the resulting mesh is smoothed, using Laplacian smoothing. The latter steps of subdivision, decimation and smoothing occurs every J iterations. The overall algorithm can be seen in Fig. 1. The overall process starting from the farthest point sampling is repeated for a set number of iteration, or until some stopping criterion is achieved.

3 Results

3.1 Short Axes Contours

In order to assess our algorithm, we performed a series of experiments. In a first instance we applied our method to a set of parallel SA contours and observed our technique’s ability to recover the ground truth mesh. We also assessed the effect the amount of SA had on the quality of the mesh by increasing the amount of SA contours provided initially from 2 to 15, between the base and the apex. Our validation consisted of calculating the distances between vertices of our generated mesh to the ground truth mesh (mean shape of the statistical model) and can be seen in Fig. 2.




Fig. 2. a) The ground truth mesh (mean shape of a statistical model) and the synthesized contours in red, b) Our resulting mesh given the red contours, c) The mesh-to-mesh distance error. The resulting mesh has been clipped at the level of the most basal contour.

3.2 Short Axes and Long Axes Contours

In a second experiment, we assessed the impact providing LA contour information had. Similarly to the first experiment, we looked at the varying the amount




Fig. 3. Impact of increasing the amount of SA has on the error distance between the generated mesh and the ground truth mesh

of LA's and the effect it had on the overall mesh result, as well as look at the influence LA positions had. Four short axes were chosen as the amount of SA slices, which was chosen empirically due to the previous experiment, and an increasing amount of long axis contours were provided as the initial input by adding a rotated version of the initial LA contour every iteration. Likewise, the same analysis was performed, but only keeping two LA contours every time and only changing their position.

3.3 Non-parallel, non-coincidental contours

Synthetic Data As this method is intended to work with clinical data, we simulated a typical set of cardiac contours that is usually obtained. From the mean shape model, 10 SA contours (resulting in 10mm inter-slice distance) as well as 2 LA contours (simulating 2 and 4 chambers view) were generated. This is the equivalent of a typical clinical segmentation set. The contours were then perturbed simulating a breathing misalignment by applying small random translations and rotations (in plane as well as out of plane). The resulting non-parallel, non coincidental contours were given as an input to our algorithm. The resulting mesh can be seen in Fig. 5(a). The distance map overlayed on the mesh is the distance from the resulting mesh to the input contours. The algorithm is also tested on a right ventricle, as seen in Fig. 5(b).

Clinical Data The algorithm was applied to real clinical data to show how our algorithm responds to both normal and pathological cases.

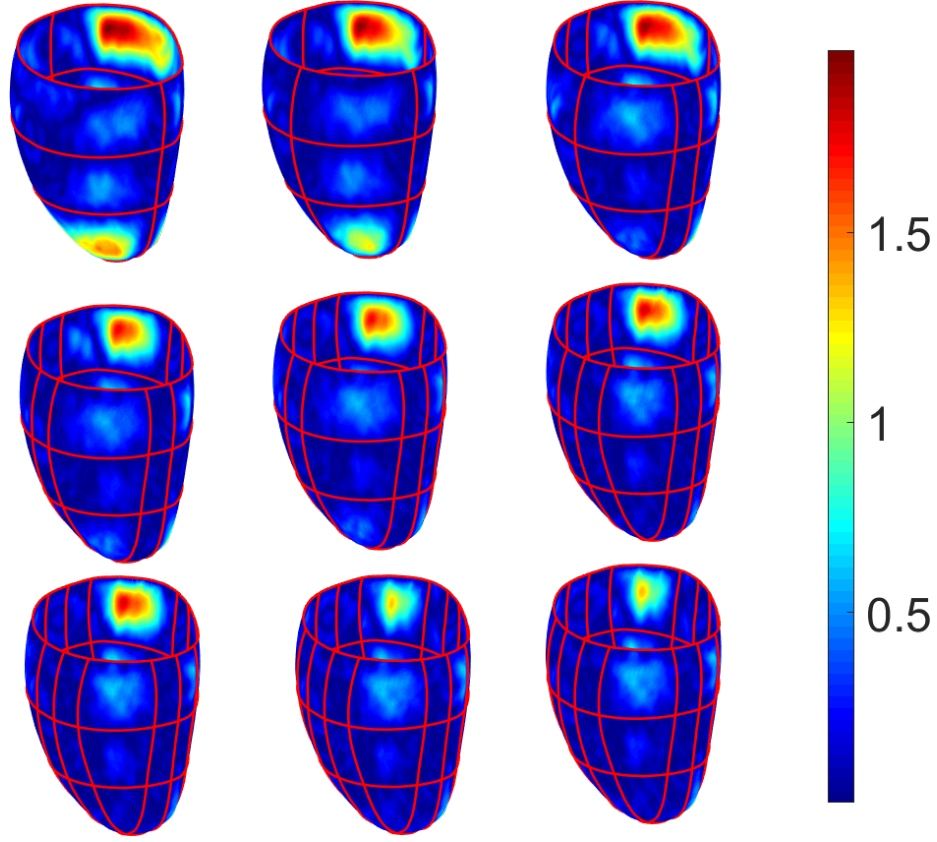


Fig. 4. Effect increasing the amount of LA contours have on the error distance between the generated mesh and the ground truth mesh.

4 Discussion

To be able to assess our algorithm, several experiments were run; ranging from synthetic data using a mean statistical shape model, to severe pathological data. The method was applied to a stack of synthesised parallel contours and the result can be seen in Fig. 2. In order to capture the discrepancy between the ground truth and the resulting mesh, vertex-to-mesh distances were calculated. Due to having no data points in the ground truth at the base (see Fig. 2(a)), the resulting mesh was clipped using the most basal contour plane. It can be seen in Fig. 2(c) that the maximum discrepancy error is around 4.5mm. This error is the distance between the two meshes most apical areas. As we do not have information at the apex due to the absence of LAs, it is expected that this is where the discrepancy would be highest. The other distances however are minimal and appear no greater than 2mm. To be able to assess the impact the amount of SA contours have on the algorithm, the distance between meshes




Fig. 5. Distance map from the resulting geometrical mesh to the simulated cardiac input data. The contours were transformed by small out of plane rotations and translations.




Fig. 6. Different clinical contours, delineated by an expert, on a severe case of hypertrophic cardiomyopathy.




Fig. 7. Resulting surface mesh including the left Endocardial and Epicardial ventricle as well as the right Epicardial ventricle obtained from a) contours for a normal case, and b) contours belonging to a severe hypertrophic cardiomyopathy patient.

was calculated by adding an extra short axis to the input data, in an iterative manner, starting with 2 SA until a stack of 15. The resulting box plot showing the distances can be seen in Fig. 3. Figure 3 shows that there is a significant reducing of errors up until 8 slices, after which the decrease in the distance error is not as notable. In this experiment, 8-9 slices represented an average clinical dataset as a SA stack is generally separated by 10mm in between. The outliers make up the error distances at the base of the mesh, where no information was provided. The impact that long axes have on the provided data was then assessed. A stack of 4 short axis slices were selected at equidistant positions from base to apex, and two long axes were placed at various different positions. Results showed that the position of the LA contours had no major contribution to minimising the error distance. We then looked at increasing the number of LA contours and evaluated the errors, which can be seen in Fig. 4. It can be seen that as the amount of LA increases, the error decreases, as well as the area with no information, which is due to the smoothness term being constrained by the data fitting terms. We then looked at generating non parallel, non coincidental contours to simulate clinical data, both with the left and right ventricle. It can be seen in Fig. 5 that where the contours were, the error distance is negligent, and the biggest error occurs in the areas where there is no information. Despite being having non parallel, non coincidental contours, the resulting meshes are visually pleasing. Finally clinical data was used, for both a normal and pathological case. The clinical contours for the pathological case (patient with hypertrophic cardiomyopathy) can be seen in Fig. 6. The resulting meshes can be observed in Fig. 7. As we have no ground truth we can only rely on the visual aspects of the resulting meshes.

5 Conclusion

We presented a tool for surface reconstruction applicable to sparse, non parallel, non coincidental, heterogeneous cardiac contours. Our technique works on the

left ventricle as well as on the right ventricle. We have empirically shown by means of the exposition of a variety of results, that our technique results in reconstructed interpolating surfaces with very good fitting to the input contours. In the areas to be filled by the interpolated surfaces, our technique presents very pleasant visual appearance. Although the absence of ground truth data precludes fair validations, our technique shows very good agreement against synthetic data.

References

1. H.J. Arevalo, F. Vadakkumpadan, E. Guallar, A. Jebb, P. Malamas, K.C. Wu, and N.A. Trayanova. Arrhythmia risk stratification of patients after myocardial infarction using personalized heart models. *Nature Communications*, 7, 2016.
2. M. Zou, M. Holloway, N. Carr, and T. Ju. Topology-constrained surface reconstruction from cross-sections. *ACM Transactions on Graphics (TOG) - Proceedings of ACM SIGGRAPH 2015*, 34, 2015.
3. K. Sunderland, W. Boyeong, C. Pinter, and Fichtinger G. Reconstruction of surfaces from planar contours through contour interpolations. *Image-Guided Procedures, Robotic Interventions, and Modeling, SPIE Proceedings*, 9415, 2015.
4. D. Deng, J. Zhang, and L. Xia. Three-dimensional mesh generation for human heart model. *Life System Modeling and Intelligent Computing. Communications in Computer and Information Science*, 98:157–162, 2010.
5. A.A. Young, B.R. Cowan, S.F. Thrupp, W.J. Hedley, and L.J. DellItalia. Left ventricular mass and volume: Fast calculation with guide-point modelling on MR images. *Radiology*, 2(216):597–602, 2000.
6. P. Lamata, S. Niederer, D. Nordsletten, D.C. Barber, I. Roy, D.R. Hose, and N. Smith. An accurate, fast and robust method to generate patient-specific cubic Hermite meshes. *Medical Image Analysis*, 15(6):801–813, 2011.
7. Xènia Albà, Marco Pereañez, Corné Hoogendoorn, Andrew J Swift, Jim M Wild, Alejandro F Frangi, and Karim Lekadir. An algorithm for the segmentation of highly abnormal hearts using a generic statistical shape model. *IEEE transactions on medical imaging*, 35(3):845–859, 2016.
8. William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3D surface construction algorithm. In *ACM siggraph computer graphics*, volume 21, pages 163–169. ACM, 1987.
9. B. Villard, E. Zacur, E. DallArmellina, and V. Grau. Correction of slice misalignment in multi-breath-hold cardiac MRI scans. *Statistical Atlases and Computational Models of the Heart. Imaging and Modelling Challenges. STACOM 2016. Lecture Notes in Computer Science (LNCS)*, 10124, 2017.
10. A. Lopez-Perez, R. Sebastian, and J.M. Ferrero. Three-dimensional cardiac computational modelling: methods, features and applications. *Biomedical Engineering Online*, 35(14), 2015.
11. Z. Zhang, K. Konno, and k. Tokuyama. 3D terrain reconstruction based on contours. *Ninth International Conference on Computer Aided Design and Computer Graphics*, pages 325–330, 2005.
12. Z. Wang, N. Geng, and Z. Zhang. Surface mesh reconstruction based on contours. *International Conference on Computational Intelligence and Software Engineering*, pages 325–330, 2009.
13. L. Liu, C. Bajaj, J.O. Deasy, and T. Ju. Surface reconstruction from non-parallel curve networks. *Eurographics*, 27, 2008.

14. W. Shi C. Minas N.G. Keenan T. Diamond G. Durighel G. Montana D. Rueckert S.A. Cook D.P. O'Regan A. de Marvao, T. Dawes. Population-based studies of myocardial hypertrophy: high resolution cardiovascular magnetic resonance atlases improve statistical power. *J Cardiovasc Magn Reson*, 16(16), 2015.
15. W. Bai, A. W. Shi, T.J.W. Dawes de Marvao, D.P. O'Regan, S.A. Cook, and D. 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.
16. Karl Rohr, H Siegfried Stiehl, Rainer Sprengel, Thorsten M Buzug, Jürgen Weese, and MH Kuhn. Landmark-based elastic registration using approximating thin-plate splines. *IEEE Transactions on medical imaging*, 20(6):526–534, 2001.
17. R. Sprengel, K. Rohr, and H.S. Stiehl. Thin-plate spline approximation for image registration. *Engineering in Medicine and Biology Society*, 18, 1996.
18. Michael Garland and Paul S Heckbert. Surface simplification using quadric error metrics. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 209–216. ACM Press/Addison-Wesley Publishing Co., 1997.