

Non-manifold Mesh Extraction from Time-varying Segmented Volumes used for Modeling a Human Heart

Martin Bertram, Gerd Reis, Rolf H. van Lengen, Sascha Köhn, and Hans Hagen

Center for Artificial Intelligence (DFKI), Kaiserslautern, Germany

Abstract

We present a new algorithm extracting and fairing surfaces from segmented volumes composed of multiple materials. In a first pass, the material boundaries in the volume are smoothed considering signed distance functions for the individual materials. Second, we apply a marching-cubes-like contouring method providing initial meshes defining material boundaries. Non-manifold features emerge along lines where more than two materials encounter. Finally, the mesh geometry is relaxed in a constrained fairing process. We use our algorithm to construct a heart model from segmented time-varying magnetic resonance images. Information concerning the heart ontology is used to merge certain structures to functional units.

Categories and Subject Descriptors (according to ACM CCS): G.1.2 [Numerical Analysis]: Approximation of surfaces and contours I.3.5 [Computer Graphics]: Curve, surface, solid, and object representations

1. Introduction

Segmented (tagged) volumes are used to describe the geometry of multiple solid shapes composed of different materials or associated with different attributes. Volumes of this type are created by segmenting computer tomography (CT) or magnetic resonance (MR) volumes, see figure 1. Automatic segmentation methods exist for certain applications. In the case of medical applications, however, human interaction is often indispensable to recover certain anatomical structures. Manual segmentation is mostly performed on slices of a vol-

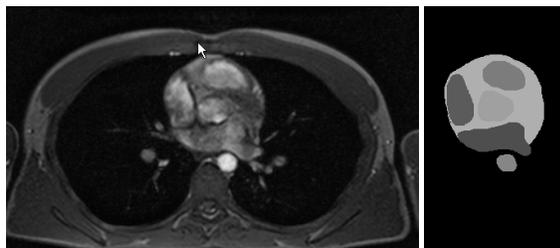


Figure 1: Magnetic resonance image (left) and segmented image (right).

ume, degrading the fairness of contours orthogonal to these slices.

In solid modeling, complex geometric shapes can be represented by signed distance functions [PT92], which are easily transformed into a segmented volume containing multiple solids associated with different materials. Hierarchical structures like octrees provide a memory-efficient representation. Segmented volumes can also be obtained from numerical simulations, like computational fluid dynamics.

Since every voxel in a segmented volume is associated with a unique attribute, material boundaries are not smooth, at all. Additional errors arising from measurement and segmentation degrade the quality of such contours. Required is a contouring method extracting and fairing topologically correct surface boundaries.

In the present work, we contribute a contouring method providing smooth meshes for every boundary between two materials. These meshes have consistent boundary features at locations where three or more materials meet. The union of all meshes represents a non-manifold surface structure. Our fairing process involves three steps:

- Fairing the underlying volume based on linear combinations of signed distance functions.

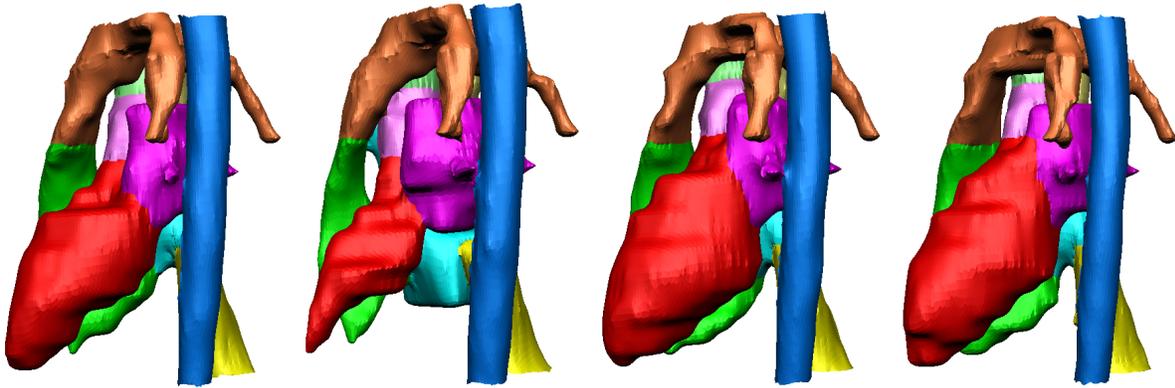


Figure 2: Heart model extracted from time-varying segmented volume, illustrating the in- and deflation of the left ventricle.

- Extraction of consistent meshes for every boundary between two materials.
- Constrained Laplacian smoothing of the extracted meshes.

In both fairing steps, the geometric error is bounded by the width of one (or a fixed number of) grid cells.

The reconstruction of a beating human heart from magnetic resonance images is a challenging problem in medical imaging, since the heart is constantly moving while the images are recorded. Triggered by echocardiogram (ECG), corresponding MR images from different heart cycles are grouped together to individual volumes. Small movements of the patient or differences in the heart cycles will lead to misaligned slices in the volume. These errors may still be present after segmentation and need to be reduced by proper fairing.

We use our algorithm for constructing the geometry of a human heart, illustrated in figure 2, based on manually segmented magnetic resonance images. Information concerning the heart ontology is used to merge certain structures to functional units, like the left ventricular complex, and to improve certain aspects of the model. Our method has been developed and is well proven for the reconstruction of anatomical structures, but it is not restricted to medical applications.

These are the contents of our work: In section 2, we summarize related work. Section 3 contains the method description, composed of the three parts: volume fairing, contouring, and constrained surface fairing. In section 4, we present our construction of a heart model and we conclude in section 5.

2. Related Work

Surface extraction from volumes (*contouring*) is typically performed by variants of the marching-cubes algorithm [LC87], where a triangulated approximation of an isosurface

is constructed for every cell in a hexahedral grid defining a three-dimensional density field. Less popular, but much simpler are contouring methods on the dual grid associated with voxel boundaries. The first approach of this kind, *elastic surface nets* [Gib98, LG99] provides quadrilateral meshes with topological ambiguities, however. Further dual methods are described by [JLSW02] and [Nie04]. An interesting extension is the reconstruction of sharp features, like ridges and cusps [KBSS01, OB02, JLSW02].

If only one material boundary is represented, a signed distance function (SDF) can be used to define an isosurface [HDD*92, PT92, CC00, JS01]. This SDF is negative inside and positive outside the material, and the boundary surface is simply the zero set. In addition to a mesh representing the surface, normal vectors can be sampled from the SDF's gradient. Adaptive volume representations for SDF's [FPRJ00, LBDM02] are not necessarily less memory efficient than surface-based approaches. Efficient contouring of hierarchical volumes, based on octrees is also available [WKE99].

The contouring problem becomes more difficult when only binary volumes [COKLY00, Whi00] are given, i.e. no continuous scalar field exists. In this case, the extracted meshes need to be smoothed, such that they still remain valid isosurfaces of the gridded volume. For example, an extracted surface can be considered as a grid composed of isocurves on three sets of slices. Constrained iterative smoothing of this curve network produces fair surfaces that are still valid isosurfaces of the binary data [NGH*03]. Other fairing methods include feature-sensitive subdivision [OB02], Laplacian smoothing [Gib98], level sets [Whi00], diffusion and low-pass filtering [DMSB99, LM99, Tau00, HG00], as well as variational methods [HHB93, WW92, SK01].

Multiple materials complicate the contouring process, since the material boundaries define a non-manifold, i.e. there exist feature lines with more than two attached sheets.

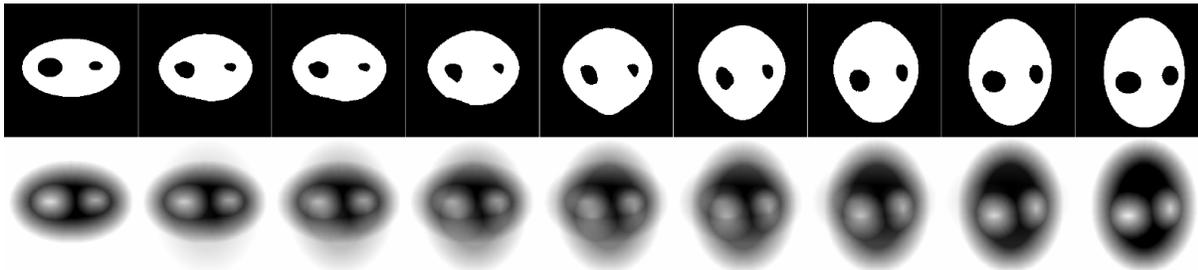


Figure 3: Interpolation between two slices (left and right upper image) based on linear combinations of signed distance functions (bottom row).

Even for tetrahedral grids a variety of topologically involved cases need to be considered [BF95]. In the case of continuous material fractions where the percentage of every material can be specified for each grid point, efficient meshing methods exist [BDS*03, GSA*03]. A generalization of marching cubes to segmented data with associated material probabilities is described by [HSSZ97] and was recently used for non-manifold extraction [AH04].

In contrast to most other approaches, our method reconstructs material boundaries from segmented (tagged) volumes, only. It is designed for medical applications, like reconstruction from MR volumes, where slices may not be well aligned causing severe artifacts. These are reduced by a volume interpolation and fairing pass based on linear combinations of signed distance functions. The same method can be used for interpolation between time steps in time-varying applications. Our mesh extraction is efficient and simple, since it does not rely on complex look up tables containing all possible topological cases. Constrained Laplacian smoothing is used for final fairing. Topological inconsistencies due to sampling and segmentation errors are eliminated by an ontology-based classification of boundary meshes.

3. Mesh Extraction and Fairing

In the following, we present the details of our method. The first step is a fairing process of the volume, reducing artifacts due to misaligned slices. Second, we extract meshes from the fair volume. These need to be smoothed, again, to avoid gridding artifacts. This last step, will not correct errors larger than the width of one grid cell, such that slicing artifacts need to be reduced by preprocessing the volume.

3.1. Volume Fairing

In the case of volumes assembled from individual 2D images, like MR data, the distances of these slices are greater than the width of one pixel (in our application, the image resolution is $\epsilon_p \approx 1mm$ opposed to $\epsilon_s \approx 7mm$ slice distance).

Since the slices assembled to a volume are not recorded simultaneously (in our case they even correspond to different heart cycles), these may not fit exactly together. To reduce errors of this kind, material boundaries in the volume need to be smoothed across the slices in a first pass.

For volume fairing, we rely on signed distance functions SDF's. This method has been used by Jones/Chen [JC94] for reconstructing material boundaries from a set of contours by extracting the zerosurface of a 3D SDF. To gain efficiency, our approach computes SDF's only on the individual slices. In addition, we use this technique for interpolating and fairing boundaries of multiple materials rather than only interpolating boundaries of two materials. In cases where errors of measurement are not introduced by certain slices, the method can be applied successively for fairing in the three canonical directions.

We first consider the problem of interpolation between two binary slices, i.e. using only two different material attributes, as illustrated in figure 3. We compute an Euclidean signed distance function (ESDF) for the material boundary in each image. Therefore, the boundary pixels are initialized with zero distance. Starting at this border the pixel front is propagated iteratively, thereby filling the respective distance values $(1, \sqrt{2}, 2, \sqrt{5}, \dots)$. Since the distances are filled incrementally there is no need for a special treatment of self-intersections. In a second step the distance map is transformed into a signed distance map simply by flipping the sign inside the feature domain, resulting in negative values inside and positive values outside the feature.

After constructing the ESDF's $s_0(x, y)$ and $s_1(x, y)$ for the two reference images, additional images can be obtained by linear interpolation with respect to $t \in [0, 1]$:

$$s_t(x, y) = (1-t) s_0(x, y) + t s_1(x, y). \quad (1)$$

The corresponding image is obtained by filling all pixels where $s_t(x, y)$ is negative with material, see figure 3.

Considering the ESDF's of three adjacent slices s_{i-1} , s_i , and s_{i+1} , fairing is obtained by replacing the middle slice by

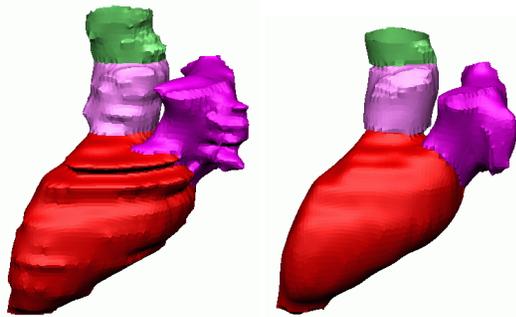


Figure 4: The effect of volume smoothing on the cavity complex of the left ventricle.

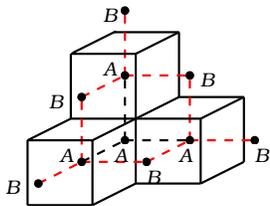


Figure 5: For every edge connecting voxels with two different materials, here denoted by A and B, a quadrilateral is generated.

the material distribution derived from

$$\bar{s}_i = (1 - 2\alpha) s_i + \alpha (s_{i-1} + s_{i+1}), \quad (2)$$

Where $\alpha \in [0, \frac{1}{3}]$ is a constant trading off smoothness versus accuracy. We note that this fairing process can shift the material boundary by at most one slice distance ϵ_s .

Both interpolation and fairing is based on convex combinations of ESDF's. In the case of multiple materials, we need to construct an ESDF $s_{m,i}$ for each material m and each participating image i . After computing the convex combinations \bar{s}_m for every material, the individual pixels of the interpolated slice are filled with the material of minimal distance,

$$\bar{m}(x,y) = \arg \min_m \{\bar{s}_m(x,y)\}. \quad (3)$$

To save memory, we store the blended ESDF of one material in a Z-buffer. The other materials are successively processed, and the Z-buffer is updated at each pixel where the Z-buffer entry is greater than the current ESDF. The corresponding material index is simultaneously stored in a frame buffer.

Problems of this interpolation / fairing technique occur when features to be connected do not overlap in the participating images. In our application, the slices are sufficiently close to connect these features. For more sophisticated matching of topological features, we refer to adaptive matching techniques [MSS92].

Gaps between adjacent materials may occur, if these ma-

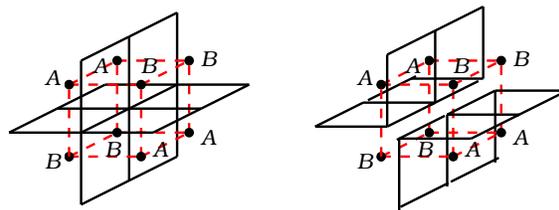


Figure 6: Solving topological ambiguities by virtually moving the boundary towards the material of lower priority and by splitting the involved vertices.

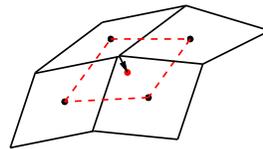


Figure 7: Two-pass smoothing strategy.

terials are not present in all participating images of the convex combination. To avoid infinite values of ESDF's, the range of each ESDF is bounded by the interval $[-d_{max}, d_{max}]$, where d_{max} is the maximal expected distance of features in adjacent slices. If topological information is available regarding individual components composed of multiple materials (for example joining vessels defined by heart ontology), these components can be treated as single “dummy” material in a first pass before separating the participating materials in a second pass. This treatment eliminates feature lines between connected components.

For our application, we first smoothed the volume using $\alpha = 0.1$ and then inserted additional slices by linear interpolation, doubling the overall number of slices. The effect of volume smoothing is illustrated in figure 4.

3.2. Contouring

The contouring applied in our method is similar to the generalized marching cubes proposed by [HSSZ97] for non-manifold meshes. Our method, however does not rely on continuous material probabilities extracted from the original MRI, since in our application too many materials are involved. In the case of three materials, a lookup table needs to contain 58 topologically different triangulations [HSSZ97], opposed to 14 in the binary case. Since a cell can be composed of up to eight materials, we decided to adopt a different approach than constructing a lookup table for all topological cases. We observed that it is much simpler to use the dual grid [Gib98, Nie04] composed of voxels rather than cells for contouring.

Between every pair of materials, a boundary mesh needs

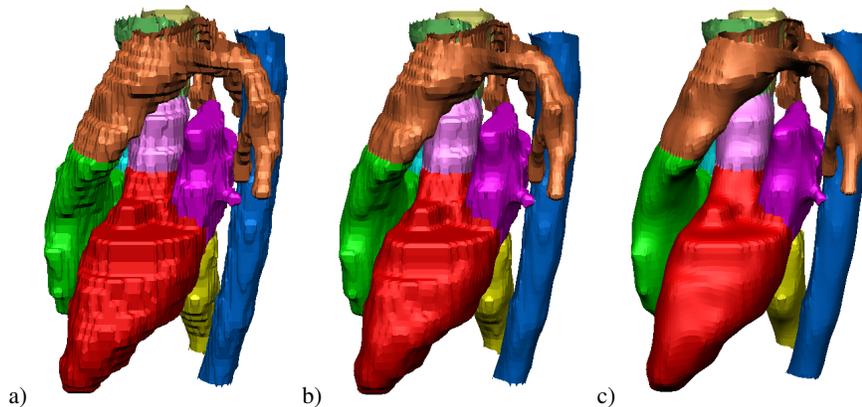


Figure 8: Meshes during the fairing process (rendered with flat shading). a) heart before fairing; b) after one constrained fairing step; c) after five fairing passes.

to be extracted. The union of all meshes defines a non-manifold surface, where only the set of meshes enclosing one particular material is guaranteed to be a manifold. The dual-grid approach assumes that the material attributes are associated with the voxel centers. This corresponds to shifting the hexahedral grid by half a cell width in each direction. The material boundaries correspond to a set of quadrilateral faces enclosing a connected complex of voxels filled with the same material, see figure 5.

For every pair of adjacent grid voxels with respect to a 6-neighborhood, we add a quadrilateral (*face*) representing their common boundary to our mesh. Every face is tagged with both material attributes. Hence, for each cell (corresponding to the grid used by MC) intersected by a surface component, one initial vertex is generated. This vertex needs to be duplicated in certain cases, for example to separate two sheets of the same material. However, such topological corrections are in general not unique, since there exists no continuous density field suggesting which material components are connected and which are separated.

We resolve this ambiguity by introducing a priority ordering of material attributes. Material components of greater priority are merged in ambiguous cases. This is implemented by moving the mesh components virtually toward the material of lower priority. Since every face is associated with two materials, the direction for this virtual movement is unique. Consequently, vertices incident to faces moving in opposite directions are split into multiple (topological) vertices, associated with the individual mesh components, see figure 6. In cases where at least one polygon contains both components of a split vertex, we do not perform the split, since we want to generate quadrilaterals, only. These cases occur in regions with three or more adjacent materials.

3.3. Constrained Fairing of Meshes

The meshes issued by our contouring method are topologically simple, since they are composed of quadrilateral faces and most vertices have valences between three and six. The geometry, however, requires some fairing. Our fairing method is inspired by Laplacian smoothing, where every vertex is replaced by the centroid of its neighbors. Since the fairing should be independent of the order in which the vertices are processed, we perform two steps: First, the centroid of every face is computed. Second, every vertex is replaced by the centroid of its incident-face centers (this is done by traversing the faces, again, to avoid the need of storing links from vertices to faces), see figure 7.

Accuracy of the smoothed mesh is enforced by the constraint that every vertex must remain in the dual-grid cell, where it has initially been created. Due to this constraint, the mesh can move at most by the width of half a grid cell. (We note that the range of movement can be augmented to allow relaxation across multiple voxels.) This fairing procedure is repeated multiple (say five) times. Spikes in the smoothed surface are avoided by dropping the constraint in the last pass, allowing the vertices to leave their originating voxels. Again, the error is bounded, since all incident faces are located in the 27-neighborhood of the dual-grid voxel.

Our fairing process is illustrated in figure 8. We observe that feature lines along common boundaries of three or more materials are still visible, for example the curve across the pulmonary artery (top of figure 8) where the boundary of the heart domain (not visible) is attached. If necessary, such features can be smoothed by assigning priorities to the individual surface components, such that a vertex attached to more than two components is relaxed with respect to the union of the two adjacent meshes of highest-priority.

We note that our fairing method does not conserve material volume. If this is required, correction operators (in-

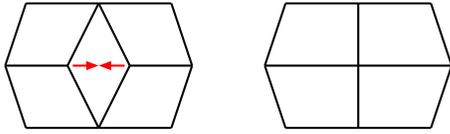


Figure 9: Regularization of mesh structure by collapsing quadrilaterals with two opposing valence-three vertices.

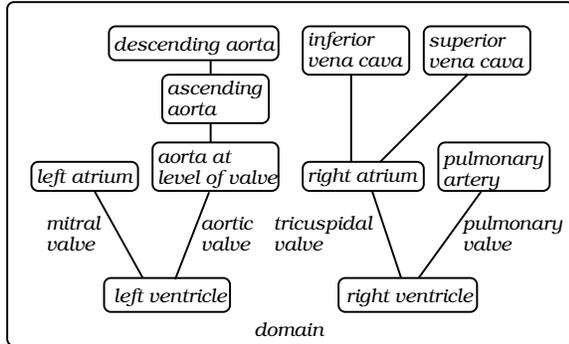


Figure 10: Heart ontology defining left and right ventricular complexes embedded in the heart domain.

volving multiple vertices) need to be introduced, in order to restore the individual material volumes after each vertex manipulation. A non-shrinking fairing approach is described by [Tau00].

When analyzing the meshes issued by contouring and fairing, we observed a great number of valence-three vertices in surfaces diagonal to the voxel grid. Mesh regularity can be improved by an additional pass collapsing all faces with two diagonally adverse valence-three vertices. These vertices are simply re-joined and their common face is eliminated, see figure 9.

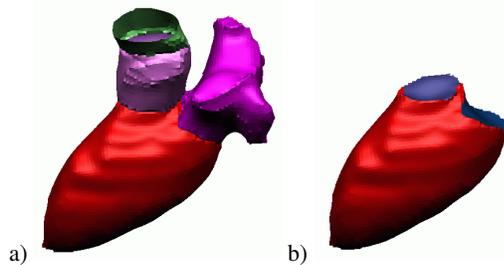


Figure 11: a) Left ventricle with atrium (right) and aorta; b) left ventricle with mitral and aortic valves.

4. Constructing a Heart Model

Considering the construction of a human heart model from time-varying segmented MR volumes, there are some application-specific aspects to mention. In the following, we summarize topological corrections and numerical results for our heart model.

4.1. Validating Model Ontology

The algorithm described above provides one surface component for each pair of materials that have a common boundary (each component may be composed of multiple parts). Considering the components initially extracted from our segmented heart volume, we are not able to decide whether two adjacent materials are connected or not. For example, if two vessels that are not physically connected touch each other due to proximity, they will produce a common surface component. The correct solution would duplicate such surface components and merge them with the remaining surfaces of both individual vessels. Topological errors of this kind are often overlooked during the segmentation, which is performed on single slices.

The only way to decide which meshes need to be connected and to eliminate “topological waste” is by taking into account an ontology of heart anatomy, summarized in figure 10. For example, we know that blood coming from the lungs passes the left atrium entering the left ventricle through the mitral valve. From the left ventricle, it is pumped through the aortic valve and the aorta into the body. Hence, all surface components involving the left ventricle, except for the mitral and aortic valves, can be grouped to one single component. Analogously, we process the left atrium and the ascending aorta, which was subdivided into an upper and a lower segment, see figure 11.

After processing the left ventricular complex, we consider the right ventricle. From heart ontology, we know that the superior and the inferior vena cava is connected to the right atrium, where the blood enters the heart after traversing the body. From there, it enters the right ventricle through the tricuspidal valve and is pumped via the pulmonary artery into the lungs. Again, we can group all meshes bounding the right ventricle, except for the tricuspidal and pulmonary valves, together and proceed with the other structures in the right ventricular complex, leaving only their inner boundaries as separate components. We note that potential boundary meshes between left and right complex are now duplicated and represented independently in both structures.

The process of identifying surface parts of individual functional units and combining them to a model consistent with heart ontology is illustrated in figure 12. We observe, for example, that inner structures like the right atrium touch the heart boundary and thus produce two surface components, one with the heart domain and one with the outer ma-

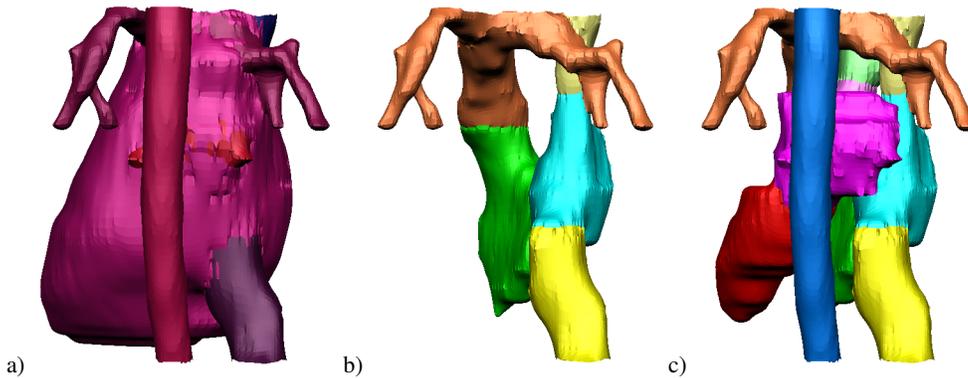


Figure 12: Combining meshes associated with functional units. a) All initially extracted meshes; b) right ventricular complex with pulmonary artery; c) left and right ventricular complexes combined.

resolution	no. faces	no. vertices	extraction time	fairing time
$36 \times 152 \times 256$	48,821	47,952	0.142 sec	0.185 sec
$71 \times 304 \times 512$	182,711	179,790	0.895 sec	0.697 sec

Table 1: Mesh sizes and computation times for contouring and mesh fairing of the first time step.

terial. The combined meshes obtained for inner structures are depicted in figure 12 c).

4.2. Numerical Results

The initial MRI data set has 36 slices in a resolution of 304×512 pixels for each of 22 time steps. Triggered by the r-wave of the echocardiogram (ECG), a sequence of slices was recorded. Slices from different heart cycles were grouped together, according to their timely distance from the r-wave. The ECG was not part of the data set.

Due to the dynamics of the beating heart, registration techniques compensating a patient’s movement during the recording process, for example by identifying common landmarks in consecutive slices, were not available for the acquisition of this data set. This implies that some of the slices combined to a single time step may not fit exactly, since they were recorded within different cycles where the heart may behave slightly different. This lack of consistency makes the reconstruction problem even more challenging. We observed that only a combination of ESDF-based volume smoothing and mesh fairing provides acceptable results.

All computations were performed on a consumer-grade PC equipped with a 2GHz Processor and 512MB RAM. Our software implementation of the volume fairing approach used about 30 seconds for computing and fairing the ESDF’s of one time step. We note that a GPU-based implementation of this part would be feasible, due to its similarity with the z-buffer method. The computation times for mesh extraction

and fairing (based on five constrained and one unconstrained Laplacian smoothing steps) are much lower, since after an initial sweep only cells involved with the surface need to be processed. Table 1 shows the results for extracting and smoothing meshes for one time step at two different resolutions (for interactive rendering, we used the meshes of lower resolution).

We found that 22 time steps were sufficient for interactive animation of our heart model. Additional time steps can be generated by volume interpolation, as described in section 3.1. The color plate shows screen shots of our interactive renderer using a clipping plane to visualize the inner structures of the heart.

5. Conclusions

We presented an efficient and robust mesh extraction method for segmented volumes encompassing multiple materials. Despite of topological complexity and local geometric errors of the initial volumes, our method produces highly accurate and smooth meshes. This was achieved by the combination of two different fairing approaches processing the underlying volume and the extracted meshes. The volume fairing requires that structures to be connected overlap in adjacent slices. For both fairing methods, the movement of surface components is bounded. Our algorithm was successfully used to reconstruct and visualize a time-varying human heart from a set of segmented magnetic resonance images. Future work will be directed at the interactive manipulation

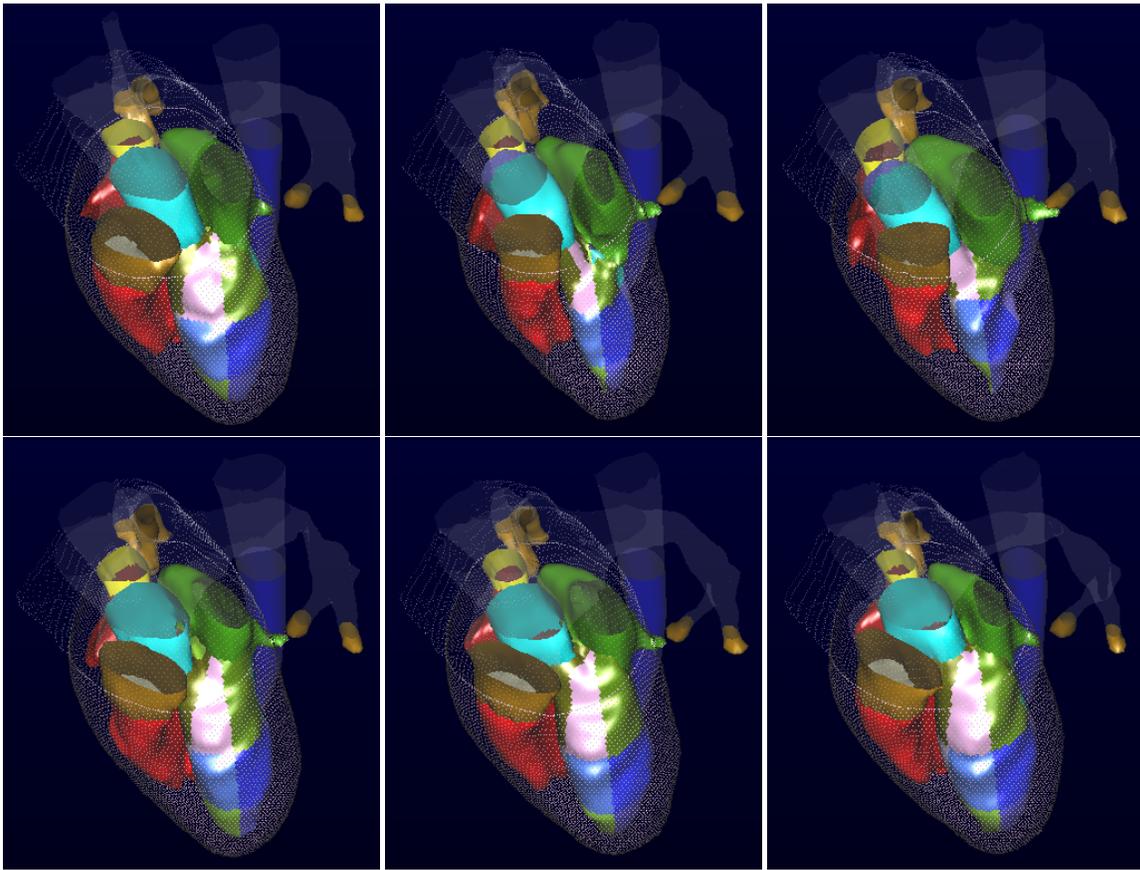
and validation of the extracted meshes with respect to additional data, like echocardiography.

Acknowledgements

This work was supported by the German Federal Ministry of Education and Research (BMBF), under contract NR 01 IW A02, through project VES. The authors want to thank Daniel Steffen for his work on the segmentation and Karsten Schwenk for implementing an interactive surface viewer.

References

- [AH04] ADOLFSSON J., HELGESSON J.: Generation of smooth non-manifold surfaces from segmented image data. *Masters Thesis, Zuse Institute Berlin (ZIB) and Lund Institute of Technology* (2004). 3
- [BDS*03] BONNELL K. S., DUCHAINEAU M. A., SCHIKORE D. A., HAMANN B., JOY K. I.: Material interface reconstruction. *Transactions on Visualization and Computer Graphics* 9, 4 (2003), 500–511. 3
- [BF95] BLOOMENTHAL J., FERGUSON K.: Polygonization of non-manifold implicit surfaces. *Siggraph* (1995), 309–316. 3
- [CC00] COSTA L. F., CESAR R. M.: *Shape Analysis and Classification: Theory and Practice*. Series on Image Processing. CRC Press LCC, 2000. 2
- [COKLY00] COHEN-OR D., KADOSH A., LEVIN D., YAGEL R.: Smooth boundary surfaces from binary 3d datasets. In *Volume Graphics* (2000), Springer, pp. 71–78. 2
- [DMSB99] DESBRUN M., MEYER M., SCHROEDER P., BARR A.: Implicit fairing of irregular meshes using diffusion and curvature flow. *Siggraph* (1999), 317–324. 2
- [FPRJ00] FRISKEN S. F., PERRY R. N., ROCKWOOD A. P., JONES T. R.: Adaptively sampled distance fields: a general representation of shape for computer graphics. *Siggraph* (2000), 249–254. 2
- [Gib98] GIBSON S. F.: Constrained elastic surface nets. *Medical Image Computation and Computer Assisted Interventions* (1998), 888–898. 2, 4
- [GSA*03] GREGORSKI B. F., SIGETI D. E., AMBROSIANO J. J., GRAHAM G., WOLINSKY M., DUCHAINEAU M. A., HAMANN B., JOY K. I.: Multiresolution representation of datasets with material interfaces. In *Hierarchical and Geometrical Methods in Scientific Visualization* (2003), Springer Heidelberg, pp. 99–117. 3
- [HDD*92] HOPPE H., DE ROSE T., DUCHAMP T., McDONALD J., STUETZLE W.: Surface reconstruction from unorganized points. *Siggraph* (1992), 71–78. 2
- [HG00] HUBERLI A., GROSS M.: Fairing of non-manifolds for visualization. *Visualization* (2000), 407–414 & 581. 2
- [HHB93] HAGEN H., HAHMANN S., BONNEAU G. P.: Variational surface design and surface interrogation. *Eurographics* (1993), 447–459. 2
- [HSSZ97] HEGE H. C., SEEBASS M., STALLING D., ZÖCKLER M.: A generalized marching cubes algorithm based on non-binary classifications. *Technical Report SC 97-05, Zuse Institute Berlin (ZIB)* (1997). 3, 4
- [JC94] JONES M. W., CHEN M.: A new approach to the construction of surfaces from contour data. *Computer Graphics Forum* 13, 3 (1994), 75–84. 3
- [JLSW02] JU T., LOSASSO F., SCHAEFER S., WARREN J.: Dual contouring of hermite data. *Siggraph* (2002), 339–346. 2
- [JS01] JONES M., SATHERLEY R.: Shape representation using space filled sub-voxel distance fields. *Shape Modeling and Applications* (2001), 316–325. 2
- [KBSS01] KOBBELT L. P., BOTSCH M., SCHWANECKE U., SEIDEL H.-P.: Feature-sensitive surface extraction from volume data. *Siggraph* (2001), 57–66. 2
- [LBDM02] LANEY D. E., BERTRAM M., DUCHAINEAU M. A., MAX N.: Multiresolution distance volumes for progressive surface compression. *3D Data Processing Visualization and Transmission* (2002), 470–479. 2
- [LC87] LORENSEN W. E., CLINE H. E.: Marching cubes: a high resolution 3d surface construction algorithm. *Siggraph* (1987), 163–169. 2
- [LG99] LEVENTON M. E., GIBSON S. F.: Models generation from multiple volumes using constrained elastic surface nets. In *Information Processing in Medical Imaging* (1999), Springer, pp. 388–393. 2
- [LM99] LEVY B., MALLET J. L.: Discrete smooth interpolation: Constrained discrete fairing for arbitrary meshes. *Technical Report* (1999). 2
- [MSS92] MEYERS D., SKINNER S., SLOAN K.: Surfaces from contours. *ACM Transactions on Graphics* 11, 3 (1992), 228–258. 4
- [NGH*03] NIELSON G. M., GRAF G., HOLMES R., HUANG A., PHIELIPP M.: Shrouds: optimal separating surfaces for enumerated volumes. *Vis.Sym* (2003), 75–84 & 287. 2
- [Nie04] NIELSON G. M.: Dual marching cubes. *IEEE Visualization* (2004), 489–496. 2, 4
- [OB02] OHTAKE Y., BELYAEV A. G.: Dual/primal mesh optimization for polygonized implicit surfaces. *ACM Solid Modeling* (2002), 171–178. 2
- [PT92] PAYNE B., TOGA A.: Distance field manipulation of surface models. *Computer Graphics and Applications* 12, 1 (1992), 65–71. 1, 2
- [SK01] SCHNEIDER R., KOBBELT L.: Geometric fairing of irregular meshes for freeform surface design. *Computer Aided Geometric Design* 18, 4 (2001), 359–379. 2
- [Tau00] TAUBIN G.: Geometric signal processing on polygonal meshes. *Eurographics'00 – State of the Art Report* (2000). 2, 6
- [Whi00] WHITAKER R. T.: Reducing aliasing artifacts in isosurfaces of binary volumes. *IEEE Symposium on Volume visualization* (2000), 23–32. 2
- [WKE99] WESTERMANN R., KOBBELT L., ERTL T.: Real-time exploration of regular volume data by adaptive reconstruction of isosurfaces. *The Visual Computer* 15, 2 (1999), 100–111. 2
- [WW92] WELCH W., WITKIN A.: Variational surface modeling. *Siggraph* (1992), 157–166. 2



Color Plate: Interactive rendering showing one heart cycle.