High-Quality Multi-Tissue Mesh Generation for Finite Element Analysis

Panagiotis Foteinos[1,2] and Nikos Chrisochoides[2]

[1] Computer Science Department, College of William and Mary, Virginia, USA, pfot@cs.wm.edu
[2] Computer Science Department, Old Dominion University, Virginia, USA, nikos@cs.odu.edu

Abstract. Mesh generation on 3D segmented images is a fundamental step for the construction of realistic biomechanical models. Mesh elements with low or large dihedral angles are undesirable, since they are known to underpin the speed and accuracy of the subsequent finite element analysis. In this paper, we present an algorithm for meshing 3D multi-label images. A notable feature of our method is its ability to produce tetrahedra with very good dihedral angles respecting, at the same time, the interfaces created by two or more adjoining tissues. Our method employs a Delaunay refinement scheme orchestrated by special point rejection strategies which remove poorly shaped elements without deteriorating the representation of the objects' anatomical boundaries. Experimental evaluation on CT and MRI atlases have shown that our algorithm produces watertight meshes consisting of elements of very good quality (all the dihedral angles were between 19 and 150 degrees) which makes our method suitable for finite element simulations.

1 Introduction

Meshing multi-labelled medical images (like those obtained by segmenting MRI or CT images) provides the means for constructing accurate bio-mechanical models for subsequent finite element analysis. Multi-material mesh generation imposes challenges, since it should meet two conflicting requirements: *fidelity* and *quality*.

Fidelity measures the capability of the mesher to preserve the boundaries formed by two or more adjoining tissues. Quality regards the shape of the elements: tetrahedra with small or large dihedral angles (i.e., low quality tetrahedra) result in interpolation errors and in ill-conditioned stiffness matrices undermining in this way the accuracy and speed of the associated finite element analysis [1].

The difficulty in mesh generation is that the need to preserve high-curvature creases of the object's surface (i.e., high fidelity) deteriorates the quality of the meshes; on the other hand, the quality of mesh elements should be as high as possible when dealing with isotropic materials [2].

In this paper, we propose a Delaunay meshing algorithm able to respect the interfaces of multi-material domains and produce tetrahedra with very good dihedral angles and radius-edge ratios (and therefore very good aspect ratios), offering at the same time control over the size of the mesh.

1.1 Previous work

In the literature, there has been work on multi-tissue meshing but the issue of high quality has not been adequately addressed.

Meyer *et al.* [3] employ a particle-based scheme producing watertight meshes that respect the interfaces formed by adjoining tissues. However, elements with practically zero dihedral angles (slivers) do appear in the final meshes. Furthermore, the execution times reported range from 3 to 12 hours even for small datasets.

Liu *et al.* [4] compress a body-centered cubic lattice (BCC) using a point-based registration method. The dihedral angles, however, can be as low as 4°. Also, the uniform lattice results in an unnecessary large number of elements in the interior of the objects.

Chentanez *et al.* [5] model the insertion of needles into soft tissues. The resulting conforming meshes are observed to consist of elements of angles more than 10.3° and less than 160°. It is worth noting that their goal is to represent a 1-dimensional curvilinear object (the needle) as a subset of a single-tissue mesh, which is a goal quite different from ours.

Goksel and Salcudean [6] present a variational meshing technique which combines both meshing and segmentation. They report minimum angles as large as 20°. The synthetic data they used for the evaluation is a sphere, that is, a 2-manifold. Usually, multi-tissue domains consists of complicated geometries, i.e., non-manifold parts which intersect with more than one tissues. These domains impose challenges to any meshing technique and are the focus of this work.

Zhang *et al.* [7] develop an octree-based meshing algorithm. Although edge-contraction and smoothing schemes are employed for quality improvement, the authors do not report the dihedral angles observed in their meshes.

Hu *et al.* [8] and Hartmann and Kruggel [9] develop uniform meshes for multi-material domains achieving dihedral angles more than 10°, albeit without good fidelity: their meshes suffer from the “staircase” effect.

Based on previous work on single material Delaunay surface [10] and volume meshing [11], Pons *et al.* [12] present a meshing algorithm for multi-tissue domains. Recently, Boltcheva *et al.* [13] extend the work of Pons *et al.* [12], so that 0- and 1-junctions are preserved in the final meshes. Both these methods apply *sliver exudation* [14] in order to improve the quality of the mesh. Edelsbrunner and Guoy [15], however, have shown that in most cases sliver exudation does not remove all poor tetrahedra: elements with dihedral angles less than 5° survive. Indeed, Pons *et al.* [12] and Boltcheva *et al.* [13] report dihedral angles as low as 4°.

1.2 Our contribution

In this paper, we present a Delaunay refinement algorithm for meshing multi-tissue medical data so that the boundaries between neighboring tissues are conforming. It works directly on segmented images meshing both the surface and the volume of the tissues.

A notable feature of our method is its ability to produce tetrahedra with very good dihedral angles: in all the experiments on synthetic and real images we ran, our algorithm produces watertight meshes consisting of tetrahedra with dihedral angles larger than 19° and smaller than 150°.

The technique we employ for quality improvement is inspired by the work of Shewchuk [16]. Therein, poor tetrahedra are eliminated by inserting the center of their circumball, giving priority to tetrahedra with larger radius-edge ratio. Shewchuk, however, meshes input domains bounded by polyhedral surfaces. In this paper, we extend this technique to deal with multi-tissue domains bounded by curved surfaces. The main difficulty is that vertices near the surface might be inserted during quality improvement. This fact in turn hurts fidelity: edges that cross interfaces or holes appear. To overcome this problem, we propose special *point rejection strategies*. They improve the quality of elements preventing the insertion of points near the surface; rather, carefully chosen points are inserted precisely on the surface. This allows to achieve both good quality and good fidelity meshes.

The rest of the paper is organized as follows: Section 2 outlines the concept of Delaunay refinement. The multi-tissue capability of our algorithm and the point rejection strategies are described in Section 3. Lastly, Section 4 presents results on CT and MR multi-label images and Section 5 concludes the paper.

2 Background

Delaunay meshes have been shown to successfully approximate the surface of both manifold and non-manifold surfaces [10], due to the properties of the *restricted Delaunay triangulations*, first introduced by Amenta and Bern [17].

Let $V \subset \mathcal{R}^3$ be a set of vertices and $\mathcal{D}(V)$ their Delaunay triangulation. Any Delaunay triangulation satisfies the *empty ball* property: the circumscribing open ball (a.k.a *circumball*) of each tetrahedron in $\mathcal{D}(V)$ does not contain any vertex.

The *voronoi point* of a tetrahedron $t \in \mathcal{D}(V)$ is defined as the center (a.k.a *circumcenter*) of t's circumball. The *voronoi edge* of a triangle $f \in \mathcal{D}(V)$ is the segment containing those points of $\mathcal{R}^3$ such that (a) they are equidistant from f's vertices and (b) they are closer to f's vertices than to any other vertex in V.

Let $\mathcal{O}$ be the multi-label domain to be meshed. We denote $\mathcal{O}$'s surface with $\partial\mathcal{O}$. The *restriction* of $\mathcal{D}(V)$ to $\mathcal{O}$ (denoted with $\mathcal{D}_{|\mathcal{O}}(V)$) is defined as the set of tetrahedra in the triangulation whose circumcenter lies inside $\mathcal{O}$.

It can be shown [10, 11] that if V samples $\partial\mathcal{O}$ sufficiently densely, then the set of boundary triangles (a.k.a *restricted facets*) of $\mathcal{D}_{|\mathcal{O}}(V)$ is a good approximation of $\partial\mathcal{O}$ in a both topological and geometric sense. The approximation guarantees

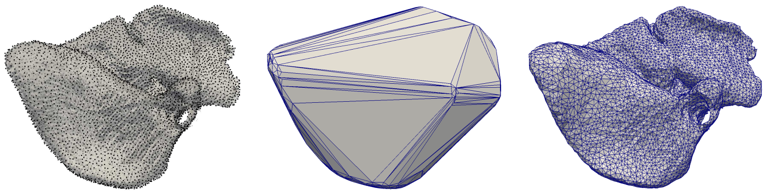

Fig. 1. (a) Sample set V of a liver's surface. (b) The Delaunay triangulation $\mathcal{D}(V)$ of the samples. (c) The restricted triangulation $\mathcal{D}_{|\mathcal{O}}(V)$.

hold as long as $\partial\mathcal{O}$ does not have sharp corners. This is a reasonable assumption, since biological tissues do not exhibit sharp features on their surface. See Figure 1 for a single-tissue example. The same idea extends to more than one tissues as well.

As an interesting consequence of the way $\mathcal{D}_{|\mathcal{O}}(V)$ is defined, only the voronoi edges of the restricted facets intersect the surface $\partial\mathcal{O}$, a property that we will exploit in Section 3 to improve quality.

3 Our method

The input of our algorithm is an image $\mathcal{I}$ containing the multi-material object $\mathcal{O}$. Image $\mathcal{I}$ can be seen as a function $f : \mathcal{R}^3 \mapsto \{0, 1, 2, \ldots, n\}$, such that $f(p)$ is the label that point $p \in \mathcal{R}^3$ belongs to. More precisely, $f(p)$ is the label of the voxel that p lies in. Usually, a label of 0 denotes voxels outside $\mathcal{O}$.

Points on the surface $\partial\mathcal{O}$ of object $\mathcal{O}$ are classified as those points lying in a voxel of label i which is incident to at least one other voxel of label j, such that $i < j$. In this way, surface $\partial\mathcal{O}$ contains not only the portions of the image that separate $\mathcal{O}$ from the background, but it also contains the interfaces that separate any adjoining tissues. The goal is to recover $\partial\mathcal{O}$ and mesh the volume (induced by $\partial\mathcal{O}$) at the same time.

Our algorithm first creates a box by inserting its 8 corners. The box contains $\mathcal{O}$ such that the (shortest) distance between the box and $\partial\mathcal{O}$ is larger than $2\delta\sqrt{2}$. Parameter δ is the only parameter that the users have to specify. This parameter determines how densely $\partial\mathcal{O}$ will be sampled: lower values indicate a denser sampling which in turn implies a better surface approximation. Notice that the calculation of the corners of the box is a quite trivial task, since it requires just one image traversal.

Next, the Delaunay triangulation of these corners is computed. This triangulation is the initial mesh (consisting of 12 tetrahedra) where the actual refinement starts from.

The refinement is governed by 2 steps, namely, *mesh conformity* and *point rejection quality improvement*. Upon termination, the tetrahedra whose circum-

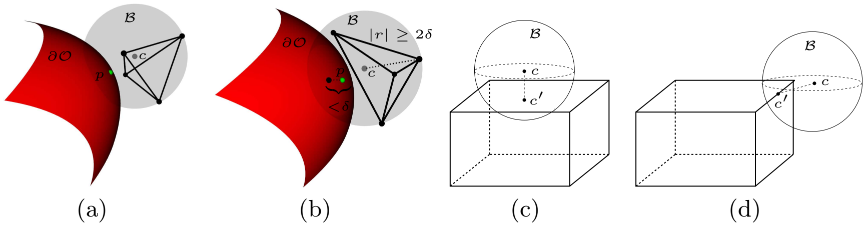

Fig. 2. (a) The closest surface point p to circumcenter c is inserted, (b) c is inserted but p is not, (c)-(d) c does not lie inside the box and therefore it is not inserted. Its projection c' is inserted and the vertices closer than δ to c' are deleted from the mesh.

center belongs to label i constitute the mesh representing the i^{th} tissue. Below, we outline each step separately.

3.1 Mesh conformity

As noted in Section 2, vertices on $\partial\mathcal{O}$ have to be inserted in order for the mesh boundary (i.e., triangles incident to 2 or more tetrahedra of different labels) to be a good approximation of $\partial\mathcal{O}$. For this reason, we keep track of the tetrahedra whose circumball $\mathcal{B}$ intersects the surface $\partial\mathcal{O}$. We call such elements *intersecting* tetrahedra.

Suppose that an intersecting tetrahedron t is found. We compute the closest surface point —say p —to the center c of t's circumball $\mathcal{B}$. To facilitate the computation of such a point, we make use of an image euclidean distance transformation [18]. If p is not closer than δ to any other surface vertex (already inserted in the mesh), then p is inserted (see Figure 2(a)). Otherwise, and if the radius of $\mathcal{B}$ is larger than 2δ, c is inserted instead (see Figure 2(b)). In this way, we can show that this step does not cause the insertion of infinite number of vertices and therefore, termination is not compromised.

For the same reason, we also require that no vertex is ever inserted outside the box. When the circumcenter c of an intersecting tetrahedron is chosen for insertion, however, c might lie outside the box. To prevent such cases, c is rejected and its projection on the box is inserted instead. See Figure 2(c) and Figure 2(d) for a couple of examples.

At the end of this step, all the vertices that do not lie on $\partial\mathcal{O}$ are deleted from the triangulation. At this moment, the restricted facets of the mesh are a good approximation of $\partial\mathcal{O}$, because the vertices remained in the triangulation form a dense sample of $\partial\mathcal{O}$ (see Section 2). Also, we can show that no 2 vertices are closer than δ and this is why δ controls the size of the mesh.

3.2 Point rejection quality improvement

Our algorithm keeps track of poor tetrahedra, i.e., tetrahedra with small or large dihedral angles. Poor tetrahedra are eliminated by inserting their circumcenter.

Priority is given to the tetrahedra with higher radius-edge ratio as in [16]. The radius-edge ratio of a tetrahedron t is defined as the length of t's circumball radius divided by the length of t's shortest edge.

Problems arise, however, when the circumcenter of a poor tetrahedron (about to be eliminated) lies close to the surface. If this is the case, the restricted facets in the triangulation are not any more a good approximation of $\partial\mathcal{O}$. See Figure 3(a) for an example: the boundary facets have vertices that do not lie precisely on the surface.

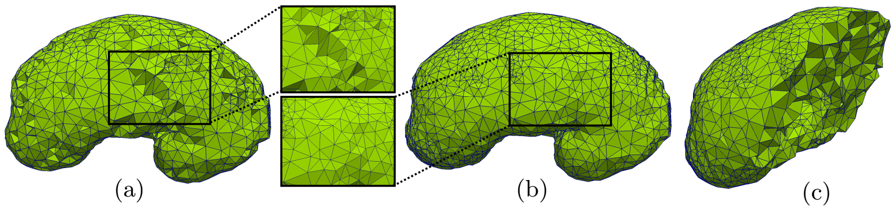

Fig. 3. Meshes for a kidney. All dihedral angles are between 19° and 150°. (a) No extra care has been taken to preserve fidelity and holes appear. (b) The point rejection strategies prevented the creation of holes. (c) A cross section of the mesh in (b).

To overcome this issue, we propose special *point rejection strategies*. Their goal is to make sure that all poor tetrahedra are eliminated without inserting points close to the surface.

Our algorithm first tries to convert *illegal* facets to *legal* ones. We define legal facets to be those restricted facets whose thee vertices lie precisely on $\partial\mathcal{O}$. Conversely, a restricted facet with at least one vertex not lying on $\partial\mathcal{O}$ is called an illegal facet.

Let t be an illegal facet and e its voronoi edge (see Figure 4(a) for an illustration). Recall that e has to intersect $\partial\mathcal{O}$ (see Section 2) at a point p. Any vertex v of t which does not lie precisely on $\partial\mathcal{O}$ is deleted from the triangulation, while point p is inserted. Note that since only non-surface vertices are deleted from the triangulation and since p is inserted on $\partial\mathcal{O}$, this step does not introduce an infinite loop: points that are inserted are never deleted.

In addition, the algorithm tries to keep in the Delaunay triangulation as many legal facets as possible. Let c be the circumcenter of a poor tetrahedron considered for insertion. If the insertion of c eliminates a legal facet t (see Figure 4(b)), then c is not inserted. Instead, a point p on the intersection of $\partial\mathcal{O}$ and t's voronoi edge e is inserted.

Figure 3(b) and Figure 3(c) show how our algorithm meshed a kidney; observe that now the boundary facets have vertices lying precisely on $\partial\mathcal{O}$. In the next section, we will demonstrate that our point rejection strategies work also very well on multi-material domains.

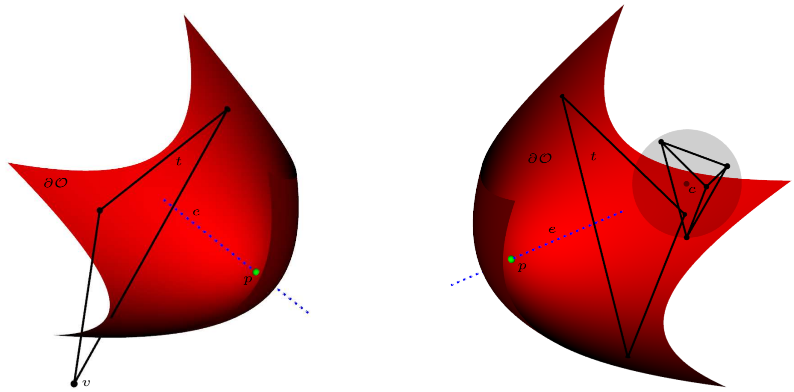

(a) (b)

Fig. 4. The point rejection strategies. (a) t is an illegal facet. (b) t is a legal facet.

4 Results

We ran our experiments on a 64 bit machine equipped with a 2.80GHz quad-core Intel i7 processor and 8GB of memory. Our algorithm was built on top of the *Computational Geometry Algorithms Library* (CGAL, http://www.cgal.org). We used the *Insight Toolkit* (ITK, http://www.itk.org) for image processing. Lastly, the *Visualization Toolkit* (VTK, http://www.vtk.org) rendered the meshes.

Figure 5(a) illustrates the output mesh obtained for a segmented CT image taken from IRCAD (http://www.ircad.fr). Similarly, Figure 5(b) depicts the output mesh obtained for the MR brodmann atlas (http://www.sph.sc.edu/comd/rorden/mricro.html). Observe that the mesh elements are of excellent quality. Although we do not give guarantees on the minimum and maximum angles achieved by our method, we observed that the point rejection strategies are able to remove elements with angles less than $19°$ and more than $150°$ (in any image input we tried) without creating an edge smaller than $\frac{\delta}{10}$. It would be interesting to theoretically investigate why elements of worse quality are eliminated so easily without introducing very small edges. We leave that exploration as a future work. See the columns "ircad" and "brodmann" of Table 1 for some statistical results.

The last row of Table 1 shows the largest tetrahedron aspect ratio. Aspect ratio is defined as the ratio of a tetrahedron's circumradius to its inradius. The reported aspect ratio is normalized such that the best aspect ratio equals 1. Therefore, the aspect ratio ranges from 1 to $+\infty$. A high aspect ratio is an indication of bad quality.

Lastly, we show that the size of the mesh can be controlled directly by parameter δ. For this reason, we ran our mesher on the same CT image (obtained by IRCAD), but this time we set the value of δ at 8. That is, we double the value of δ used to obtain the mesh of Figure 5(a). See Figure 5(c) for an illustration

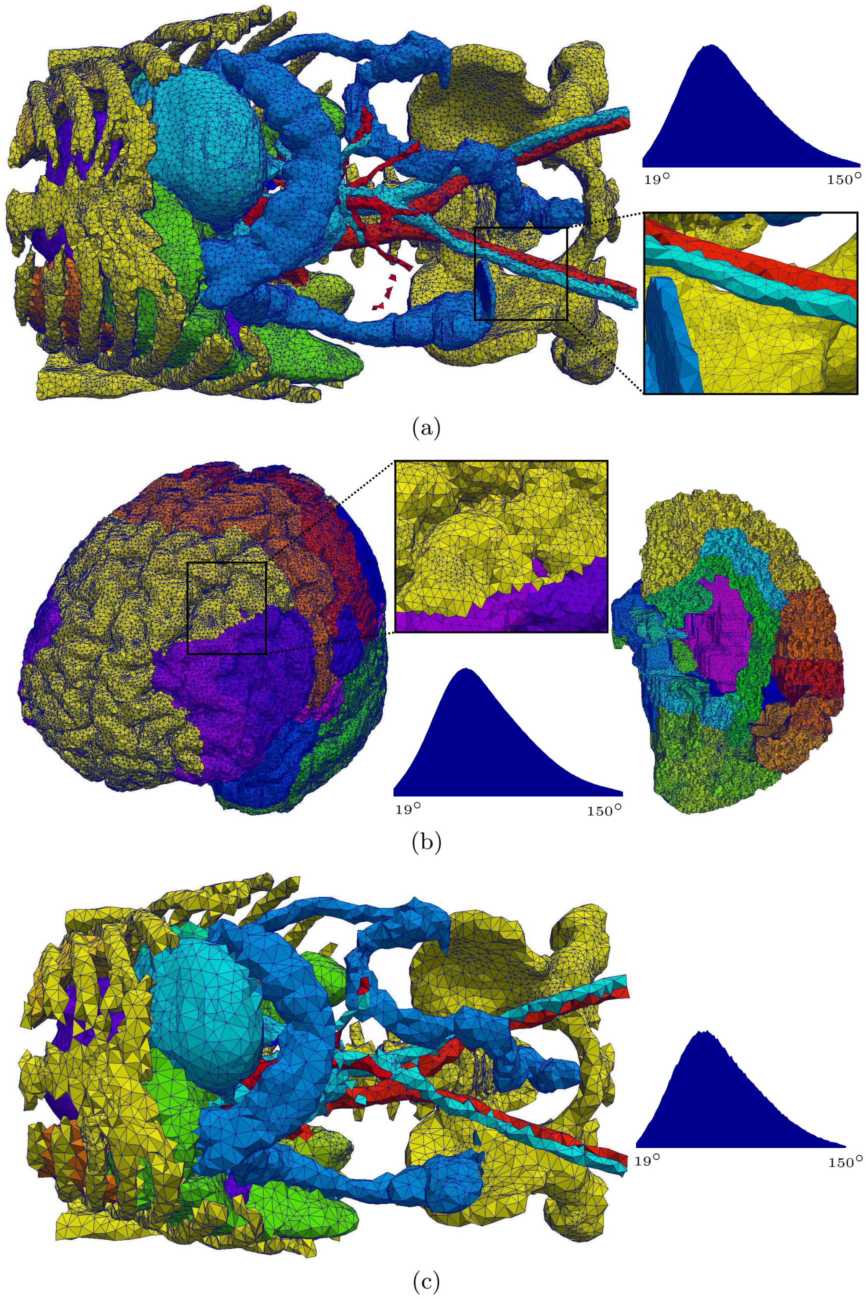

Fig. 5. Whole meshes, zoomed views, cross sections, and distributions of the dihedral angles for (a) the CT multi-label image and (b) the MR brain atlas. In (c), we show a coarser mesh on the same input image that was used in (a).

Table 1. Information about the images used for the evaluation, the chosen value for parameter δ, and some quantitative results of the final meshes produced by our algorithm.

Experiment	ircad	brodmann	ircad(coarse)
Image size	$512 \times 512 \times 219$	$181 \times 217 \times 181$	$512 \times 512 \times 219$
Image resolution (mm)	$0.961 \times 0.961 \times 2.4$	$1 \times 1 \times 1$	$0.961 \times 0.961 \times 2.4$
#Labels	20	41	20
δ (mm)	4	2	8
Time (sec)	421	$1,066$	96
#Vertices	$139,740$	$473,994$	$41,097$
#Tetrahedra	$783,445$	$2,575,220$	$173,575$
Dihedral angles (degrees)	$19-150$	$19-150$	$19-150$
Max. aspect ratio (normalized)	4.67	6.22	4.55

and the last column of Table 1 for some statistical results. Observe that the number of elements, the number of vertices, and the execution time are greatly reduced, in the expense of worse fidelity. This is an expected trade-off: the fewer elements a mesh has, the less likely it is to represent complex surface creases accurately.

To evaluate our method, we compare it with CGAL (http://www.cgal.org). A comparison with other popular meshing techniques like Tetgen (http://tetgen.berlios.de/) or Netgen (http://www.hpfem.jku.at/netgen/) is omitted in this paper, because they do not operate directly on images. Rather, they require that the surface is already meshed as a piecewise linear complex. In contrast, both our algorithm and CGAL mesh the surface and the volume at the same time.

We run CGAL on the ircad CT image and report the achieved quality. We set CGAL's sizing parameters to values that gave output meshes with size similar to the size of our mesh depicted in Figure 5(a). Furthermore, we set the quality parameters to their best theoretical values as described in [11]. Quantitative results for CGAL's output mesh are shown in Table 2. Compare it with the first column of Table 1. Observe that the quality of the CGAL mesh is lower than ours in terms of dihedral angles and aspect ratios. Another popular quality metric is the minimum scaled Jacobian value [19, 20]. It ranges from -1 to 1 with 1 being the best value. A negative value means that some elements are inverted. Both our algorithm and CGAL report positive scaled Jacobian values. In fact, the minimum Jacobian value achieved by our algorithm is 30 times larger than that achieved by CGAL. In terms of absolute numbers, we feel that the Jacobian values of our method is low, an issue we are looking into as future work.

5 Conclusions and future work

In conclusion, we have shown that Delaunay refinement techniques are able to mesh multi-material domains with tetrahedra of very good angles, which makes our method suitable for subsequent finite element analysis. The point rejection strategies proposed in this work maintain mesh conformity and high quality offering, at the same time, control over the mesh size.

Table 2. Quantitative results of the final mesh produced by CGAL.

Experiment	ircad
#Vertices	156, 902
#Tetrahedra	756, 462
Dihedral angles (degrees)	$3 - 174$
Max. aspect ratio (normalized)	16, 823

Note that surface patches of high curvature need to be meshed with more elements than patches that are not sharp. In its current state, our method meshes the surfaces uniformly. In the future, we plan to extend our method to produce graded triangular surfaces and, therefore, smaller meshes.

Acknowledgments. We thank the anonymous referees for their comments and insight that helped improve the presentation and completeness of the paper. This work is supported in part by NSF grants: CCF-1139864, CCF-1136538, and CSI-1136536 and by the John Simon Guggenheim Foundation and the Richard T. Cheng Endowment.

References

1. Shewchuk, J.R.: What is a Good Linear Element? - Interpolation, Conditioning, and Quality Measures. In: Proceedings of the 11$^{\text{th}}$ International Meshing Roundtable, Sandia National Laboratories (2002) 115–126
2. Goksel, O., Salcudean, S.E.: High-quality model generation for finite element simulation of tissue deformation. In: 12$^{\text{th}}$ International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI). MICCAI '09, Berlin, Heidelberg, Springer-Verlag (2009) 248–256
3. Meyer, M., Whitaker, R., Kirby, R.M., Ledergerber, C., Pfister, H.: Particle-based sampling and meshing of surfaces in multimaterial volumes. IEEE Transactions on Visualization and Computer Graphics **14(6)** (2008) 1372–1379
4. Liu, Y., Foteinos, P., Chernikov, A., Chrisochoides, N.: Multi-tissue mesh generation for brain images. In: Proceedings of the 19$^{\text{th}}$ International Meshing Roundtable, Chattanooga, Tennessee, USA (2010) To appear.
5. Chentanez, N., Alterovitz, R., Ritchie, D., Cho, L., Hauser, K.K., Goldberg, K., Shewchuk, J.R., O'Brien, J.F.: Interactive simulation of surgical needle insertion and steering. In: Proceedings of ACM SIGGRAPH 2009. (2009) 88:1–10
6. Goksel, O., Salcudean, S.E.: Image-based variational meshing. IEEE Transactions on Medical Imaging **30** (2011) 11–21
7. Zhang, Y., Hughes, T.J., Bajaj, C.L.: An automatic 3d mesh generation method for domains with multiple materials. Computer Methods in Applied Mechanics and Engineering **199** (2010) 405 – 415 Computational Geometry and Analysis.
8. Hu, P., Chen, H., Wu, W., Heng, P.A.: Multi-tissue tetrahedral mesh generation from medical images. In: International Conference on Bioinformatics and Biomedical Engineering (iCBBE). (2010) 1–4

9. Hartmann, U., Kruggel, F.: A Fast Algorithm for Generating Large Tetrahedral 3D Finite Element Meshes from Magnetic Resonance Tomograms. In: Proceedings of the IEEE Workshop on Biomedical Image Analysis. WBIA, Washington, DC, USA, IEEE Computer Society (1998) 184–192
10. Boissonnat, J.D., Oudot, S.: Provably good sampling and meshing of surfaces. Graphical Models **67** (2005) 405–451
11. Oudot, S., Rineau, L., Yvinec, M.: Meshing volumes bounded by smooth surfaces. In: Proceedings of the International Meshing Roundtable, San Diego, California, USA, Springer-Verlag (2005) 203–219
12. Pons, J.P., Ségonne, F., Boissonnat, J.D., Rineau, L., Yvinec, M., Keriven, R.: High-Quality Consistent Meshing of Multi-label Datasets. In: Information Processing in Medical Imaging. (2007) 198–210
13. Dobrina Boltcheva, Mariette Yvinec, Jean-Daniel Boissonnat: Mesh Generation from 3D Multi-material Images. In: Medical Image Computing and Computer-Assisted Intervention, Springer (2009) 283–290
14. Cheng, S.W., Dey, T.K., Edelsbrunner, H., Facello, M.A., Teng, S.H.: Sliver exudation. Journal of the ACM **47** (2000) 883–904
15. Edelsbrunner, H., Guoy, D.: An experimental study of sliver exudation. Engineering with Computers **18** (2002) 229–240
16. Shewchuk, J.R.: Tetrahedral mesh generation by delaunay refinement. In: Proceedings of the 14^{th} ACM Symposium on Computational Geometry, Minneapolis, MN (1998) 86–95
17. Amenta, N., Bern, M.: Surface reconstruction by Voronoi filtering. In: SCG '98: Proceedings of the fourteenth annual symposium on Computational geometry, New York, NY, USA, ACM (1998) 39–48
18. Danielsson, P.E.: Euclidean Distance Mapping. Computer Graphics and Image Processing **14** (1980) 227–248
19. Knupp, P.M.: Achieving finite element mesh quality via optimization of the jacobian matrix norm and associated quantities. Part II-A framework for volume mesh optimization and the condition number of the jacobian matrix. International Journal for Numerical Methods in Engineering **48** (2000) 1165–1185
20. Shepherd, J.F., Tuttle, C.J., Silva, C.T., Zhang, Y.: Quality improvement and feature capture in hexahedral meshes. SCI institute technical report, University of Utah (2006)