

Mesh Generation from 3D Multi-material Images^{*}

Dobrina Boltcheva, Mariette Yvinec, and Jean-Daniel Boissonnat

GEOMETRICA – INRIA Sophia Antipolis, France

Abstract. The problem of generating realistic computer models of objects represented by 3D segmented images is important in many biomedical applications. Labelled 3D images impose particular challenges for meshing algorithms because multi-material junctions form features such as surface patches, edges and corners which need to be preserved into the output mesh. In this paper, we propose a feature preserving Delaunay refinement algorithm which can be used to generate high-quality tetrahedral meshes from segmented images. The idea is to explicitly sample corners and edges from the input image and to constrain the Delaunay refinement algorithm to preserve these features in addition to the surface patches. Our experimental results on segmented medical images have shown that, within a few seconds, the algorithm outputs a tetrahedral mesh in which each material is represented as a consistent submesh without gaps and overlaps. The optimization property of the Delaunay triangulation makes these meshes suitable for the purpose of realistic visualization or finite element simulations.

1 Introduction

Motivation. Advanced medical techniques frequently require geometric representations of human organs rather than grey-level MRI or CT-scan images. The precondition for extracting geometry from a medical image is usually to *segment* it into multiple regions of interest (materials). This paper focuses on the next step: the automatic generation of meshes from 3D labelled images. These meshes are either surface meshes approximating the boundaries of anatomical structures or volume meshes of these objects.

The generation of geometric models from segmented images presents many challenges. In particular, the algorithm must handle arbitrary topology and multiple junctions. Multi-material junctions form *surface patches* (2-junctions), *edges* (1-junctions) and *corners* (0-junctions) which are expected to be accurately represented in the output mesh. However, it is still very challenging to construct a good mesh where each feature (surface patch, edge or corner) is accurately represented.

^{*} This research was supported by the FOCUS K3D Coordination Action, EU Contract ICT-2007.4.2 contract number 214993.

In this paper, we present a Delaunay refinement algorithm which addresses this issue. In a matter of seconds, the algorithm outputs a high-quality tetrahedral mesh where each anatomical structure is represented by a submesh with conforming boundaries.

Related work. There are only a few meshing strategies which directly provide volume meshes from multi-material images. Early grid-based methods deal with tetrahedral elements that are created from the original rectilinear volume subdivision [1,2]. More recently, an octree-based method has been proposed [3] that outputs tetrahedral or hexahedral meshes with consistent multi-material junctions. However, like the other grid-based methods, this algorithm does not offer a mesh quality control and elements of poor quality are always generated along boundaries. Quality improvement techniques are usually used as a post-processing step to deliver acceptable meshes.

Another powerful strategy based on Delaunay refinement, has been proposed for meshing smooth surfaces [4] and volumes bounded by smooth and piecewise smooth surfaces [5,6]. The refinement process is controlled by highly customizable quality and size criteria on triangular facets and on tetrahedra. Pons et al. [7] have adapted this mesh generation strategy to labelled images and have shown that the sizing criteria can be tissue-dependent leading to high-quality volume meshes of the different materials. However, in this work, 1- and 0-junctions are not explicitly handled and they are not accurately represented in the output mesh. It is well recognized that the presence of edges where different surface patches meet forming small angles, jeopardize the termination of Delaunay refinement in 3D [8]. A recent method [9] deals with this problem using the idea of protecting small angle regions with balls so that during the refinement no point is inserted into these protecting balls.

Very recently, Mayer et al. [10] proposed a sampling strategy for labelled images based on a dynamic particle system which explicitly samples corners, edges, and surface patches. The resulting set of well-distributed points is adapted to the underlying geometry of multi-material junctions. From these surface points, Delaunay-based meshing scheme outputs high-quality surface meshes. However, this sampling approach relies on heavy pre-processing computations and takes between 3 and 12 hours for small medical images.

Our contribution. In this paper, we propose a feature (edge and corner) preserving extension of the Delaunay refinement algorithm introduced by Pons et al. [7]. The basic idea is to explicitly sample 1- and 0-junctions before launching the Delaunay refinement and to constrain the latter to preserve these features. The algorithm first extracts, from the input 3D image, all multi-material junctions where 3 or more materials meet. Then it samples and protects these junctions with protecting balls as in [9]. The Delaunay refinement is then run using a weighted Delaunay triangulation, with the protecting balls as initial set of data points. This allows to preserve 0-junctions (corners) and to reconstruct accurately the 1-junctions (edges) with a sequence of mesh edges. In contrast to Mayer et al. [10], we sample only 0- and 1-junctions and not the 2-junctions

which are meshed by the Delaunay refinement algorithm according to user-given quality and sizing criteria. Our multi-material junction detection and sampling algorithm is very fast and has little influence on the computation time of the Delaunay refinement algorithm.

The remainder of this paper is organized as follows: Section 2 gives a brief background and recalls the Delaunay refinement algorithm. Our multi-material junction detecting and protecting strategy is presented in Section 3. Section 4 reports some results and numerical experiments which demonstrate the effectiveness of our approach for segmented medical 3D images.

2 Delaunay Refinement Mesh Generation from 3D Images

The method is related to the concept of *restricted Delaunay triangulation*, borrowed from computational geometry.




Fig. 1. (a) A set of points in the plane and its Voronoi diagram. (b) The dual Delaunay triangulation. (c) The Delaunay triangulation restricted to the blue curve is plotted with a black line. (d) The Delaunay triangulation restricted to the yellow region is composed of triangles whose circumcentres are inside the region.

Background. Let $E = \{p_0, \dots, p_n\}$ be a set of points in \mathbb{R}^3 called *sites* hereafter. The *Voronoi cell*, denoted $Vor(p_i)$, associated to the site p_i is the locus of points that are closer to p_i than to any other site in E . The *Voronoi diagram* of E , denoted $Vor(E)$, is the partition of space induced by the Voronoi cells $Vor(p_i)$. The Delaunay complex is the dual of the Voronoi diagram defined as follows: for each subset of sites $T \subset E$, the convex hull $conv(T)$ is a cell of the Delaunay complex if and only if the intersection $\bigcap_{p \in T} Vor(p)$ of the Voronoi cells of sites in T is non empty. When the set E is in general position, i.e. there are no 5 sites lying on the same sphere, the Delaunay complex is a simplicial complex called the *Delaunay triangulation* of E , denoted $Del(E)$.

Let us now consider a subset $\Omega \subset \mathbb{R}^3$ and a set of points E . We call Delaunay triangulation of E *restricted* to Ω , and denote $Del(E)|\Omega$, the subcomplex of $Del(E)$ composed of Delaunay simplices whose dual Voronoi faces intersect Ω . Fig.1 gives an example of these notions in 2D.

Let us consider a multi-material 3D image I as a function $F : \mathbb{Z}^3 \rightarrow J$, where $J = \{0, \dots, n\}$ is a set of labels where 0 represents the background and 1...n the

other materials. Each label i defines a characteristic function $f_i : \mathbb{Z}^3 \rightarrow \{0, 1\}$ such as $f_i(p) = 1$ iff $F(p) = i$ and 0 otherwise. Let $\tilde{f}_i : \mathbb{R}^3 \rightarrow \{0, 1\}$ be the trilinear interpolation of f_i . Then we define the extension $\tilde{F} : \mathbb{R}^3 \rightarrow J$ of the image function F , as follows: $\tilde{F}(p) = i$ iff $\tilde{f}_i(p) = \max_{j \in J} \{\tilde{f}_j(p)\}$. \tilde{F} defines a partition of the domain to be meshed $\Omega = \cup_{i \neq 0} \Omega_i$, where $\Omega_i = \tilde{F}^{-1}(i)$, $i \in J$. We call B_i the boundary of Ω_i and $B = \cup B_i$ denotes the locus of multi-material junctions including surface patches, edges and corners.

The Delaunay refinement meshing algorithm presented in the next section samples a set of points $E \in \mathbb{R}^3$ and builds the Delaunay triangulation of E restricted to Ω . The algorithm outputs $Del(E)|_{\Omega = \cup_{i \in J, i \neq 0} \Omega_i} Del(E)|_{\Omega_i}$ where $Del(E)|_{\Omega_i}$ is the set of tetrahedra whose circumcentres are contained in Ω_i . In other words, each tetrahedron is labelled according to the material in which its circumcentre lies. We call *boundary facets* the facets incident to two tetrahedra with different labels. These boundary facets form a watertight non-manifold surface mesh that approximates surface patches where two material meet.

Delaunay refinement algorithm. The algorithm starts by sampling a small initial set of points E on $\cup B_i$. Three points per connected component of $\cup B_i$ are sufficient. Next it calculates the Delaunay triangulation $Del(E)$ and its restrictions $Del(E)|_{\Omega_i}$ and the boundary facets. This initial approximation is then refined until it meets the following user-given quality criteria for boundary facets and tetrahedra:

- criteria for boundary facets: minimum angle α ; maximum edge length l ; maximum distance between a facet and the surface patch d ;
- criteria for tetrahedra: ratio between tetrahedron circumradius and shortest edge length β ; maximum edge length L .

Thus, the mesh refinement criteria are given by the 5-uplet (α, l, d, β, L) . A *bad element* is an element that does not fulfil criteria. *Bad facets* are removed from the mesh by inserting their *surface centres*. The *surface centre* of a boundary facet is the point of intersection between its dual Voronoi edge and a surface patch.

Bad tetrahedra are removed from the mesh by inserting their circumcentres. The algorithm inserts refinement points one by one and maintains the current set E , $Del(E)$, $Del(E)|_{\Omega}$ and boundary facets. The refinement procedure is iterated until there is no bad element left.

After the refinement, degenerated tetrahedra with small dihedral angles (slivers) are removed from the mesh using a *sliver exudation* algorithm [11].

It is proven in [6] that for appropriate choices of refinement criteria, the algorithm terminates. It delivers surface and volume meshes which form a good approximation of the image partition as soon as E is a “sufficiently dense” sample of its boundaries and volumes. However, 0- and 1-junctions are not handled explicitly and they are poorly reconstructed in the output mesh. As shown on Fig.2(1) the 1-junctions are zigzagging and the 0-junctions are not preserved.




Fig. 2. (1) Delaunay refinement 3D mesh. (2) Multimaterial junctions: five 1-junctions and two 0-junctions. (3) Sampled points on junctions. (4) Protecting balls. (5) Edge preserving Delaunay refinement 3D mesh. (6) A cut of the tetrahedral mesh. (7) Dihedral angles distribution.

3 Feature Preserving Extension

In order to constrain the Delaunay refinement algorithm to mesh properly 0- and 1-junctions, firstly we need to extract these junctions from the input image.

Multi-material junction extraction. Here, we extend the image function $F : \mathbb{Z}^3 \rightarrow J$ into a function $\tilde{F} : \mathbb{R}^3 \rightarrow J$ using the the concept of *continuous analog*, borrowed from digital geometry and topology [12,13]. Following this concept, for any point $p \in \mathbb{R}^3$, $\tilde{F}(p) = F(p_i)$, where p_i is the point of \mathbb{Z}^3 closest to p . As before, this function \tilde{F} defines a partition of the domain to be meshed $\Omega = \cup_{i \neq 0} \Omega_i$ but now $\Omega_i = \tilde{F}^{-1}(i)$, $i \in J$ is a set of cubic voxels with the same label. As before, $B = \cup B_i$ denotes the multi-material junctions which are composed of:

- 2-junctions S (surface patches composed of voxel faces) which correspond to the intersection of exactly 2 materials;
- 1-junctions L (composed of voxel edges) which are defined at the intersection of exactly 3 or 4 materials;
- 0-junctions P (defined by voxel vertices) which correspond to the intersection of 4 or more materials (at maximum 8).

As it has been stressed before, our multi-material junction extraction algorithm delivers only 0- and 1-junctions because the Delaunay refinement algorithm handles surface patches well. The result is a 1D cellular complex composed of edges $\{L_i\}$ and their end points $\{P_i\}$. Fig.2(2) shows the 1D complex obtained for the multi-material sphere which is composed of five edges and two corners.

Protecting multi-material junctions. It is well known that neighbourhoods of edges and corners are regions of potential problems for Delaunay refinement algorithm. First, as we have already seen, the usual Delaunay refinement does not reconstruct these sharp features accurately. Secondly, if we constrain these edges to appear in the mesh and if the surface patches incident to them make small dihedral angles, the usual Delaunay refinement may not terminate. However, Cheng et al.[9] have shown that if the edges are protected with *protecting balls* the Delaunay refinement terminates.

In this work, we protect 0- and 1-junctions with balls before launching the mesh refinement. We first keep all corners in $\{P_i\}$ and then we sample points on edges $\{L_i\}$ according to some user-specified density (see Fig.2(3)). This distance d between two sampled points should be at most the maximum edge length parameter for facets l . We protect each sampled point $p \in L_i$ with a ball $b = (p, r)$ where $r = 2/3*d$ (see Fig.2(4)). The protecting balls have to satisfy the following properties:

- each edge L_i is completely recovered by the protecting balls of its samples
- any two adjacent balls on a given edge L_i overlap significantly without containing each other's centres
- any two balls on different edges L_i and L_j do not intersect
- no three balls have a common intersection

After the sampling and protection step, we use the previous Delaunay refinement algorithm as follows: Each protecting ball $b = (p, r)$ is turned into a weighted point (p, r) and inserted into the initial set of points E . The Delaunay triangulation is turned into a *weighted* Delaunay triangulation where the Euclidean distance is replaced by the *weighted distance*. The weighted distance from a point $x \in \mathbb{R}^3$ to a weighted point (p, r) is defined as $\|x - p\|^2 - r^2$. All the other points inserted during the refinement are considered as points of weight zero.

The protecting balls and the weighted Delaunay triangulation guarantee two important properties. First, the segment between two *adjacent* points on any protected edge L_i remains connected with restricted Delaunay edges (see Fig.2(5)). Secondly, since the algorithm will never try to insert a refinement point into the union of protecting ball, its termination is guaranteed (see [9] for more detail).

In practice, when a multi-material 3D image is the input, the algorithm outputs high-quality tetrahedral meshes of different materials which are consistent with each other. In particular, 1-junctions are properly reconstructed with edges whose vertices lie on these junctions (see Fig.2(5)).

4 Results and Conclusion

The Delaunay refinement algorithm and its feature preserving extension have been implemented upon robust primitives to compute the Delaunay triangulation provided by the CGAL library [14].

We have tested our meshing algorithm on synthetic labelled images and on segmented medical images provided by IRCAD [15]. Figure 3 shows the meshes




Fig. 3. Meshes generated from labelled liver images: (a) A liver adjacent to a right kidney and a zoom on their edge. (b) A liver segmented into 4 anatomical regions and a zoom on one of the corners. The first row shows meshes obtained with the usual Delaunay refinement. The second row shows meshes resulting from our feature preserving extension and histograms of their dihedral angles distributions.

Table 1. Quantitative results and parameters for three different 3D images. The α in the refinement criteria is given in *degree* and l, d and L are given in *mm*. The four last rows give the computation times of different algorithm steps in seconds.

Experiment	sphere	liver-kidney	liver segments
Image size	$62 \times 62 \times 62$	$512 \times 512 \times 112$	$402 \times 356 \times 238$
Image resolution (<i>mm</i>)	$1 \times 1 \times 1$	$0.67 \times 0.67 \times 2$	$2 \times 2 \times 2$
Refinement criteria (α, l, d, β, L)	(20,10,3,4,10)	(30,12,2,4,14)	(25,14,4,4,18)
# vertices	964	6142	12381
# boundary facets	1431	5439	9646
# tetrahedra	4434	31043	64485
Junction Extraction (<i>sec</i>)	0.72	4.56	21.35
Surface meshing (<i>sec</i>)	1.74	9.99	11.04
Volume meshing (<i>sec</i>)	1.13	5.82	17.23
Sliver exudation (<i>sec</i>)	3.75	13.82	48.64

generated from two labelled liver datasets by the Delaunay refinement strategy with and without the feature preserving extension. Figure 3(a) displays the mesh of a liver adjacent to a right kidney and the 1-junction produced at the intersection of these two objects and the background. In this case, the 1-junction is a simple closed curve which has been protected with balls of radius 6mm . Figure 3(b) shows the mesh of a liver which has been cut into 4 anatomical regions. There are 8 1-junctions and 4 0-junctions produced at the intersection of these anatomical regions which have been protected with balls of radius 3mm .

Table 1 lists the quantitative results for these two liver images and the multi-material sphere on Fig.2. The refinement criteria for Delaunay refinement are given as the 5-uplet (α, l, d, β, L) defined in section 2. Note that our edge extraction and protection algorithm takes few seconds and does not penalize the Delaunay refinement which also has a reasonable computation time. A typical liver image ($512 \times 512 \times 112$) segmented into 4 different materials is meshed with target edge length of 10mm in less than half a minute.

In conclusion, we have proposed a feature preserving Delaunay refinement meshing strategy to generate conforming 3D meshes from labelled images. These meshes are targeted towards applications in the finite element methods which require conforming multi-material junctions to avoid instabilities and errors during the simulation.

References

1. Nielson, G.M., Franke, R.: Computing the separating surface for segmented data. In: VIS, pp. 229–233 (1997)
2. Hartmann, U., Kruggel, F.: A fast algorithm for generating large tetrahedral 3d finite element meshes from magnetic resonance tomograms. In: WBIA (1998)
3. Zhang, Y., Hughes, T., Bajaj, C.L.: Automatic 3d mesh generation for a domain with multiple materials. In: Meshing Roundtable, pp. 367–386 (2007)
4. Oudot, S., Rineau, L., Yvinec, M.: Meshing volumes bounded by smooth surfaces. In: Meshing Roundtable, pp. 203–219 (2005)
5. Rineau, L., Yvinec, M.: A generic software design for delaunay refinement meshing. *Comput. Geom. Theory Appl.* 38, 100–110 (2007)
6. Rineau, L., Yvinec, M.: Meshing 3d domains bounded by piecewise smooth surfaces. In: Meshing Roundtable, pp. 443–460 (2007)
7. Pons, J.-P., Ségonne, F., Boissonnat, J.-D., Rineau, L., Yvinec, M., Keriven, R.: High-quality consistent meshing of multi-label datasets. In: Karssemeijer, N., Lelieveldt, B. (eds.) IPMI 2007. LNCS, vol. 4584, pp. 198–210. Springer, Heidelberg (2007)
8. Shewchuk, J.R.: Mesh generation for domains with small angles. In: SCG, pp. 1–10. ACM, New York (2000)
9. Cheng, S.W., Dey, T.K., Ramos, E.A.: Delaunay refinement for piecewise smooth complexes. In: SODA, Philadelphia, PA, USA, pp. 1096–1105 (2007)
10. Meyer, M., Whitaker, R., Kirby, R.M., Ledergerber, C., Pfister, H.: Particle-based sampling and meshing of surfaces in multimaterial volumes. *Visualization and Computer Graphics* 14(6), 1539–1546 (2008)
11. Cheng, S.W., Dey, T.K., Edelsbrunner, H., Facello, M.A., Teng, S.H.: Sliver exudation. In: SCG, pp. 1–13 (1999)
12. Latecki, L.: 3d well-composed pictures. *Graph. Models Image Process* 59(3) (1997)
13. Françon, J., Bertrand, Y.: Topological 3d-manifolds: a statistical study of the cells. *Theoretical Computer Science*, 233–254 (2000)
14. CGAL: Computational Geometry Algorithms Library, <http://www.cgal.org>
15. IRCAD: Institut de Recherche contre les Cancers de l'Appareil Digestif, <http://www.ircad.fr>