Sampling and Meshing Curved 3D Domains by Delaunay Refinement

Jean-Daniel Boissonnat
INRIA Sophia-Antipolis

Winter School on Algorithmic Geometry
ENS-Lyon
January 2010

Applications

- visualization and graphics applications
- CAD and reverse engineering
- geometric modelling in medecine, geology, biology etc.
- autonomous exploration and mapping (SLAM)
- scientific computing : meshes for FEM

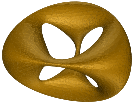

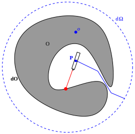

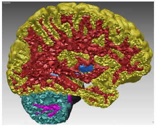

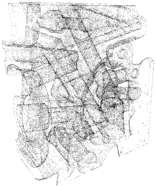

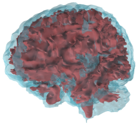

Two main issues

Sampling

- How do we choose points in the domain ?
- What information do we need to know/measure about the domain ?

Topology and Geometry

1. How do we connect the points ?
2. Under what sampling conditions can we compute a good approximation of the domain ?
3. What is a good approximation ?

State of the art : implicit surface meshing

Marching cube

Lorensen & Cline [87]
Lopez & Brodlie [03] : topological consistency
Plantiga & Vegter [04] : certified topology using interval arithmetic

Morse theory

Stander & Hart [97]
B., Cohen-Steiner & Vegter [04] : certified topology

Delaunay refinement

Ruppert [95]
Shewchuk [02]

Chew [93]
B. & Oudot [03,04]
Cheng et al. [04]

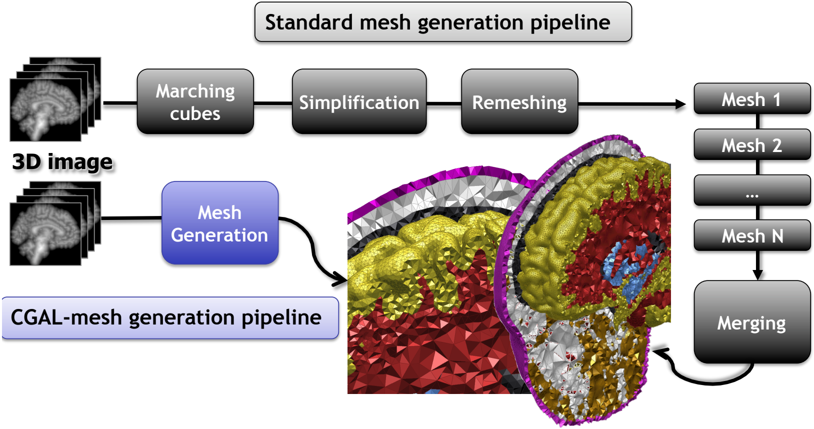
Standard mesh generation pipeline
Marching cubes
Simplification
Remeshing
Mesh 1
Mesh 2
...
Mesh N
Merging
3D image
Mesh Generation
CGAL-mesh generation pipeline

Overview

- What is a good approximation of a surface ?
- Restricted Delaunay triangulation
- Surface mesh generation
- Extensions and applications

What is a good approximation of a surface ?

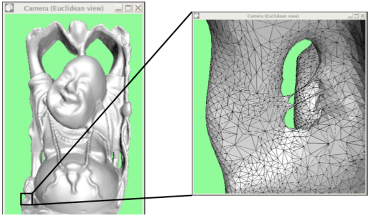

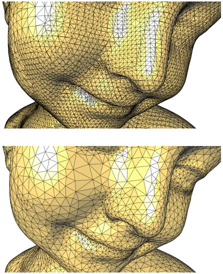

Topological equivalence

Homeomorphism

Two subsets X and Y of $\mathbb{R}^d$ are said to be homeomorphic if there exists a continuous, bijective map $f : X \rightarrow Y$ with continuous inverse f^{-1}.

Topological equivalence

Homeomorphism

Two subsets X and Y of $\mathbb{R}^d$ are said to be homeomorphic if there exists a continuous, bijective map $f : X \rightarrow Y$ with continuous inverse f^{-1}.

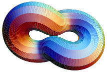

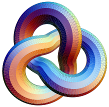

Topological equivalence

Homeomorphism

Two subsets X and Y of $\mathbb{R}^d$ are said to be homeomorphic if there exists a continuous, bijective map $f : X \to Y$ with continuous inverse f^{-1}.

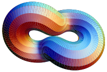

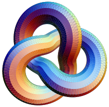

Isotopy

Two subsets X and Y of $\mathbb{R}^d$ are said to be isotopic if there exists a continuous map $f : X \times [0,1] \to \mathbb{R}^d$ such that $f(.,0)$ is the identity of X, $f(X,1) = Y$, and for each $t \in [0,1]$, $f(.,t)$ is a homeomorphism onto its image.

Distance between two sets

Hausdorff distance

$$d_H(X, Y) = \max \left(\sup_{x \in X} d(x, Y), \sup_{y \in Y} d(y, X)\right)$$

Distance between two sets

Hausdorff distance

$$d_H(X, Y) = \max\left(\sup_{x \in X} d(x, Y), \sup_{y \in Y} d(y, X)\right)$$

Fréchet distance

$$d_{\mathcal{F}}(X, Y) = \inf_h \sup_{p \in X} d(p, h(p))$$

where h ranges over all homeomorphisms from X to Y

Distance between two sets

Hausdorff distance

$$d_H(X,Y) = \max\left(\sup_{x\in X} d(x,Y), \sup_{y\in Y} d(y,X)\right)$$

Fréchet distance

$$d_{\mathcal{F}}(X,Y) = \inf_h \ \sup_{p\in X} \ d(p,h(p))$$

where h ranges over all homeomorphisms from X to Y

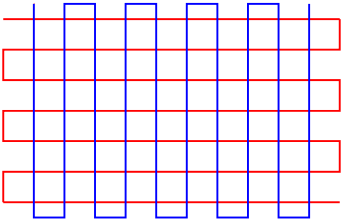

Other guarantees

- Approximation of normals
- Approximation of areas
- Approximation of curvatures
- Aspect ratio of the facets

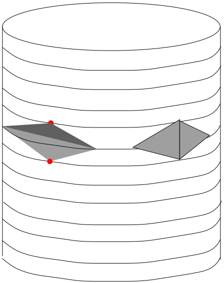

Restricted Delaunay triangulation

Data structuring by space subdivision

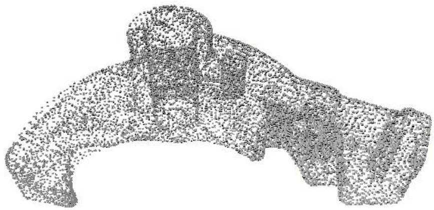

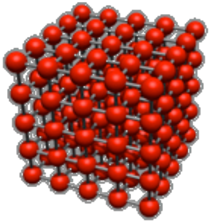

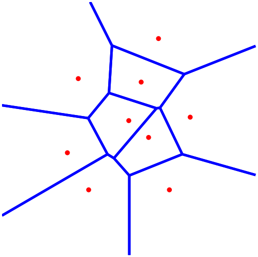

Restricted Delaunay triangulation [Chew 93]

Definition

The restricted Delaunay triangulation $\mathrm{Del}_{|\mathcal{S}}(\mathcal{P})$ is the set of simplices of the Delaunay triangulation whose dual Voronoi faces intersect $\mathcal{S}$

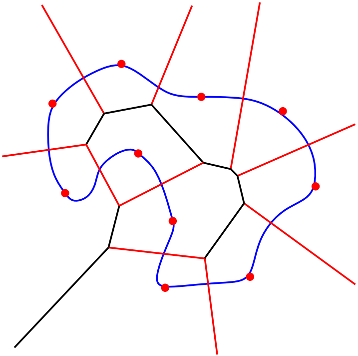

Restricted Delaunay triangulation [Chew 93]

Definition

The restricted Delaunay triangulation $\mathrm{Del}_{|\mathcal{S}}(\mathcal{P})$ is the set of simplices of the Delaunay triangulation whose dual Voronoi faces intersect $\mathcal{S}$

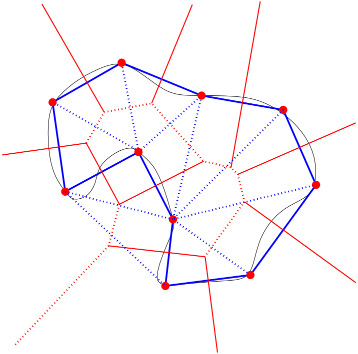

Restricted Delaunay triangulation [Chew 93]

Definition

The restricted Delaunay triangulation $\mathrm{Del}_{|\mathcal{S}}(\mathcal{P})$ is the set of simplices of the Delaunay triangulation whose dual Voronoi faces intersect $\mathcal{S}$

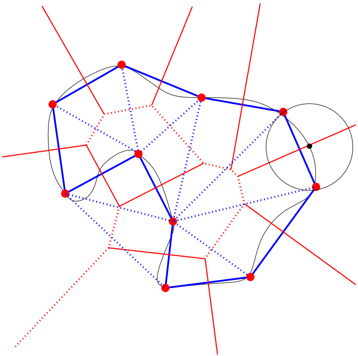

Restricted Delaunay triangulation [Chew 93]

Definition

The restricted Delaunay triangulation $\mathrm{Del}_{|O}(\mathcal{P})$ is the set of simplices of the Delaunay triangulation whose dual Voronoi faces belong to $\mathrm{Vor}_{|O}(\mathcal{P})$

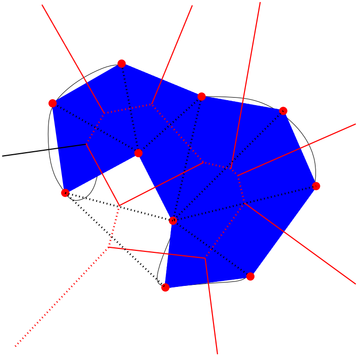

A variant of the nerve theorem

Theorem [Edelsbrunner & Shah 1997]

If S is compact and without boundary and if, for any face $f \in \mathrm{Vor}_{|S}(E)$,

1. f intersects S transversally
2. $f \cap S = \emptyset$ or is a topological ball

then $\mathrm{Del}_{|S}(E) \approx S$

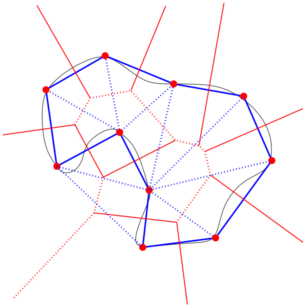

Homeomorphism

Two subsets X and Y of $\mathbb{R}^d$ are said to be homeomorphic if there exists a continuous, bijective map $f : X \to Y$ with continuous inverse f^{-1}.

A variant of the nerve theorem

Theorem [Edelsbrunner & Shah 1997]

If S is compact and without boundary and if, for any face $f \in \mathrm{Vor}_{|S}(E)$,

1. f intersects S transversally
2. $f \cap S = \emptyset$ or is a topological ball

then $\mathrm{Del}_{|S}(E) \approx S$

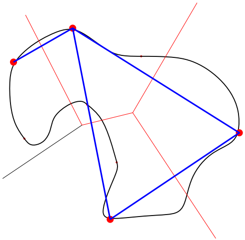

Homeomorphism

Two subsets X and Y of $\mathbb{R}^d$ are said to be homeomorphic if there exists a continuous, bijective map $f : X \to Y$ with continuous inverse f^{-1}.

Proof of the closed ball property

Barycentric subdivision

of $\text{Vor}_{\mathcal{S}}(E)$

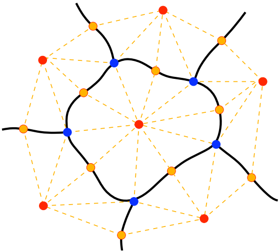

of $\text{Del}_{\mathcal{S}}(E)$

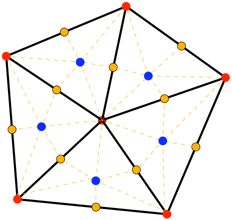

Sampling smooth surfaces

[Amenta & Bern 1998]

Local feature size

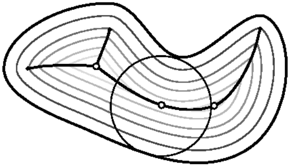

- Medial axis of $\mathcal{S}$: $M(\mathcal{S})$
  set of points with at least two closest points on $\mathcal{S}$
- Local feature size : $\mathrm{lfs}(x)$
  $\forall x \in \mathcal{S}, \;\; \mathrm{lfs}(x) = d(x, M(\mathcal{S}))$

Sampling smooth surfaces

[Amenta & Bern 1998]

Local feature size

- Medial axis of $\mathcal{S}$: $M(\mathcal{S})$
  set of points with at least two closest points on $\mathcal{S}$
- Local feature size : $\mathrm{lfs}(x)$
  $\forall x \in \mathcal{S},\ \ \mathrm{lfs}(x) = d(x, M(\mathcal{S}))$

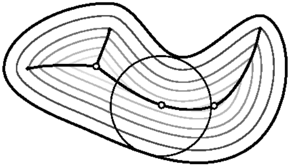

ϵ-sample of $\mathcal{S}$ (ε-covering)

$$\mathcal{P} \subset \mathcal{S}, \forall x \in \mathcal{S} = d(x, \mathcal{P}) \leq \epsilon \mathrm{lfs}(x)$$

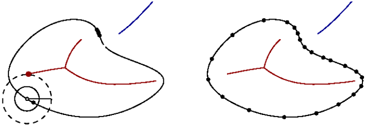

Restricted Delaunay triangulations of ε-samples

[Amenta et al. 1998-]

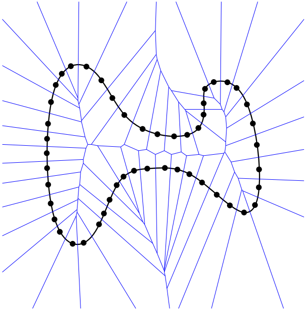

If $\mathcal{P}$ is an ε-sample of a $C^{1,1}$ surface $\mathcal{S} \subset \mathbb{R}^3$, $\varepsilon \leq 0.12$

- $\mathrm{Del}_{|S}(\mathcal{S})$ provides good estimates of
  - normals
  - areas
  - curvature [Cohen-Steiner, Morvan]
- There exists an isotopy $\phi : \mathrm{Del}_{|S}(\mathcal{P}) \to \mathcal{S}$
- $\sup_x(\|\phi(x) - x\|) = O(\varepsilon^2)$

Loose ε-samples

[B. & Oudot 2005]

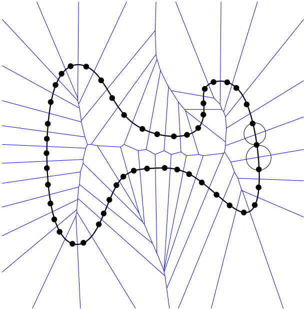

Definition

1. $\mathrm{Del}_{|S}(\mathcal{P})$ has a vertex on each connected component of S
2. for any circumscribing ball $B_f = (c_f, r_f)$ of any facet f of $\mathrm{Del}_{|S}(\mathcal{P})$, $r_f \leq \varepsilon\, \mathrm{lfs}(c_f)$

Loose ε-samples

[B. & Oudot 2005]

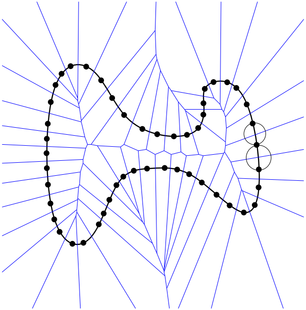

Definition

1. $\mathrm{Del}_{|S}(\mathcal{P})$ has a vertex on each connected component of S
2. for any circumscribing ball $B_f = (c_f, r_f)$ of any facet f of $\mathrm{Del}_{|S}(\mathcal{P})$, $r_f \leq \varepsilon \,\mathrm{lfs}(c_f)$

Loose ε-samples are $\varepsilon(1 + O(\varepsilon^2))$-samples

Sketch of proofs

Surfaces : local properties 1/2

Chord lemma

$$\forall p, q \in S, \|p - q\| \leq 2\varepsilon \mathrm{lfs}(p) \Rightarrow \sin(pq, T_p) \leq \varepsilon$$

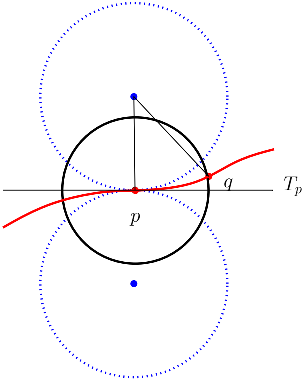

Surfaces : local properties 2/2

Facet normal lemma

Let $f = pqr$ be a facet of $\mathrm{Del}(\mathcal{P})$ and assume that $\hat{p} \geq \frac{\pi}{3}$. If the circumradius ρ_f of f is at most $\varepsilon\,\mathrm{lfs}(p)$, then $\sin(n_f, n_p) \leq 2\varepsilon$.

Proof

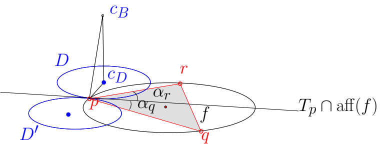

$$\sin(n_f, n_p) = \sin(pc_Bc_D) = \frac{\|p-c_D\|}{\|p-c_B\|} = \frac{\rho_D}{\mathrm{lfs}(p)}$$

$$\hat{p} = \alpha_q + \alpha_r \geq \frac{\pi}{3} \overset{\text{wlog}}{\Rightarrow} \alpha_q \geq \frac{\pi}{6}$$

$$q \notin D \cup D' \Rightarrow \|p-q\| \geq 2\rho_D \sin\alpha_q \geq \rho_D$$

$$\frac{\rho_D}{\mathrm{lfs}(p)} \leq \frac{\|p-q\|}{\mathrm{lfs}(p)} \leq \frac{2\rho_f}{\mathrm{lfs}(p)} \leq 2\varepsilon$$

Proof of homeomorphism (sketch)

The conditions of the Nerve Th. are satisfied

1. Any edge of $\mathrm{Vor}_{|S}(\mathcal{P})$ intersects S in one point

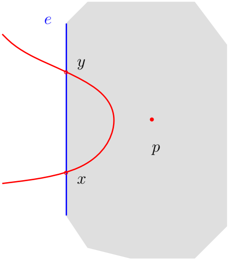

$$\|x-p\| \leq \varepsilon \,\mathrm{lfs}(x) \qquad (*)$$
$$\|y-p\| \leq \varepsilon \,\mathrm{lfs}(y) \leq \tfrac{\varepsilon}{1-\varepsilon}\,\mathrm{lfs}(x)$$
$$\Rightarrow \|x-y\| \leq 2\tfrac{\varepsilon}{1-\varepsilon}\,\mathrm{lfs}(x)$$
$$\Rightarrow [xy] \overset{|}{\sim} n_x \approx n_p$$

$[xy] \perp e^*$ and the facet lemma

$$\Rightarrow [xy] \wr n_p$$

2. Similar arguments show that faces of higher dimensions are also topological balls

$\pi : \mathrm{Del}_{|\mathcal{S}}(\mathcal{P}) \to \mathcal{S}$ is injective

$x \in \mathcal{S}$
n_x the normal to $\mathcal{S}$ at x
l_x the normal fiber $[x - r\, n_x, x + r\, n_x]$ where $r = \varepsilon\, \mathrm{lfs}(x)$

Injectivity lemma

If $\mathcal{P}$ is a loose ε-sample for $\varepsilon \leq 0.12$, then l_x intersects $\mathrm{Del}_{|\mathcal{S}}$ in at most one point

Proof of the injectivity lemma (by contradiction)

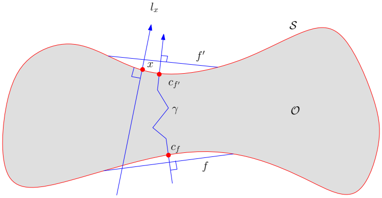

f, f' two consecutive facets of $\mathrm{Del}_{|S}(\mathcal{P})$ that intersect l_x

$T = t_1, ..., t_s$ the set of tet. intersected by l_x between f and f'

$\gamma = (c_f = c_0, c_1, ..., c_s, c_{f'} = c_{s+1})$
$c_i =$ cc of t_i, $\gamma \subset \mathrm{skel}(\mathrm{Vor}(\mathcal{P}))$

with $s_i = (c_{i+1} - c_i)/\|c_{i+1} - c_i\|$

- c_f and $c_{f'}$ are consecutive points of $l_x \cap \mathcal{S}$
  $\Rightarrow (n_{c_f} \cdot s_f) \times (n_{c_{f'}} \cdot s_{f'}) < 0$
- Facet normal lemma
  $\Rightarrow (n_{c_f} \cdot s_f) \times (n_{c_{f'}} \cdot s_{f'}) < -1 + O(\varepsilon^2)$
- Delaunay $\Rightarrow (n_x \cdot s_i) \times (n_x \cdot s_{i+1}) > 0$
  $\Rightarrow (n_x \cdot s_f) \times (n_x \cdot s_{f'}) > 0$
- Normal variation lemma
  $\Rightarrow n_{c_f} \approx n_x \approx n_{c_{f'}}$
- Contradiction !

$\pi : \mathrm{Del}_{|S}(\mathcal{P}) \to \mathcal{S}$ is surjective

If $\mathcal{P}$ is a loose ε-sample of $\mathcal{S}$ with $\varepsilon \leq 1.12$, then $\mathcal{S}$ is covered at least once by π

Proof

- any edge of $\mathrm{Del}_{|S}$ belongs to exactly two facets of $\mathrm{Del}_{|S}$
- every cc of $\mathcal{S}$ contains ≥ 1 vertex of $\mathrm{Del}_{|S}(\mathcal{P})$
- by contradiction : there exists an edge where the injectivity lemma is violated

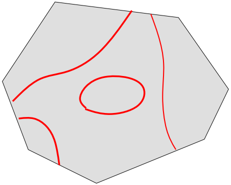

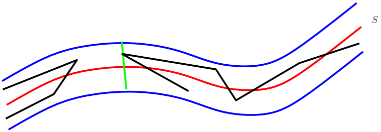

Isotopy

If $\mathcal{P}$ is an ε-sample for $\varepsilon \leq 0.12$, π induces an isotopy that maps $\mathrm{Del}_{|\mathcal{S}}(\mathcal{P})$ to $\mathcal{S}$

The isotopy moves the points by $O(\varepsilon^2)$

Isotopy

If $\mathcal{P}$ is an ε-sample for $\varepsilon \leq 0.12$, π induces an isotopy that maps $\mathrm{Del}_{|\mathcal{S}}(\mathcal{P})$ to $\mathcal{S}$

The isotopy moves the points by $O(\varepsilon^2)$

Proof

- Homeomorphism: π is bijective and bicontinuous
- Isotopy : $f : \mathrm{Del}_{|\mathcal{S}}(\mathcal{P}) \times [0,1] \rightarrow \mathcal{S}, f(x,t) = x + t\ \frac{\pi(x)-x}{\|\pi(x)-x\|}$
- Fréchet distance : trivially $\leq \varepsilon\ \sup_{x \in \mathcal{S}} \mathrm{lfs}(x)$
  for a better bound, adapt the chord lem.

Surface mesh generation by Delaunay refinement

Surface mesh generation by Delaunay refinement

$\phi : S \rightarrow \mathbb{R}$ = Lipschitz function
$\forall x \in S,\ 0 < \phi_{\text{min}} \leq \phi(x) < \varepsilon \mathrm{lfs}(x)$

[Chew 1993, B. & Oudot 2003]

ORACLE : For a facet f of $\mathrm{Del}_{|S}(\mathcal{P})$,
return c_f, r_f and $\phi(c_f)$

A facet f is bad if $r_f > \phi(c_f)$

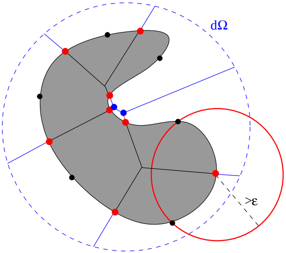

Surface mesh generation by Delaunay refinement

[Chew 1993, B. & Oudot 2003]

$\phi : S \to \mathbb{R}$ = Lipschitz function
$\quad \forall x \in S,\ 0 < \phi_{\min} \leq \phi(x) < \varepsilon \operatorname{lfs}(x)$

ORACLE : For a facet f of $\mathrm{Del}_{|S}(\mathcal{P})$,
return c_f, r_f and $\phi(c_f)$

A facet f is bad if $r_f > \phi(c_f)$

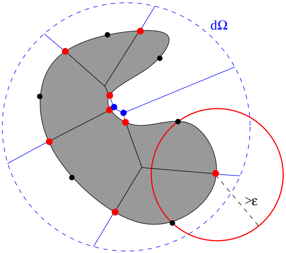

Algorithm

INIT compute an initial (small) sample $\mathcal{P}_0 \subset S$

REPEAT IF f is a bad facet
insert_in_Del3D(c_f)
update $\mathcal{P}$ and $\mathrm{Del}_{|S}(\mathcal{P})$

UNTIL all facets are good

The algorithm terminates

Properties of the output

- The output sample $\mathcal{P}$ is
  - a covering : (loose ε-sample)
    $\forall x \in \mathcal{S}, d(x, \mathcal{P}) \leq \phi(x)(1 + O(\phi^2(x))) \leq \varepsilon\,(1 + O(\varepsilon^2)\,\mathrm{lfs}(x)$
  - a packing : $\forall p \in \mathcal{P}, d(p, \mathcal{P} \setminus \{p\}) \geq \min(\phi(p), \phi(q))$
    $$\begin{aligned} &\geq \phi(p) - \|p - q\| \\ &\geq \tfrac{1}{2}\,\phi(p) \end{aligned}$$
- $|\mathcal{P}| = O\left(\int_S \frac{dx}{\phi^2(x)}\right)$
- $\mathrm{Del}_{|S}(\mathcal{P})$ is a good approximation of $\mathcal{S}$
- all facets have a bounded aspect ratio $\frac{r_f}{l_f} \leq \frac{\phi(c_f)}{\min_{x\in \mathrm{vert}(f)} \phi(x)}$

Size of the sample $= O\left(\int_S \frac{dx}{\phi^2(x)}\right)$

Proof

Let $\rho(x) = \inf\{r : |B(x,r) \bigcap \mathcal{P}| = 2\}$ and $B_p = B(p, \frac{\rho(p)}{2})$, $p \in \mathcal{P}$

$$\int_S \frac{dx}{\rho^2(x)} \geq \sum_p \int_{(B_p \cap S)} \frac{dx}{\rho^2(x)} \qquad \text{(the } B_p \text{ are disjoint)}$$

$$\geq \frac{4}{9} \sum_p \frac{\mathrm{area}(B_p \cap S)}{\rho^2(p)} \qquad \rho(x) \leq \rho(p) + \|p - x\| \leq \frac{3}{2}\, \rho(p))$$

$$\geq \frac{4}{9} \sum_p \frac{3}{16}\pi \;=\; \frac{\pi}{12}\, |\mathcal{P}|$$

$\forall x \in B_p$,

$$\rho(x) \geq \rho(p) - \|x - p\| \geq \frac{1}{2}\, \rho(p)$$
$$\rho(p) = \|p - q\| \geq \phi(p) - \|p - q\| \Rightarrow \rho(p) \geq \frac{\phi(p)}{2}$$
$$\phi(x) \leq \phi(p) + \frac{\rho(p)}{2} \leq \frac{5}{2}\, \rho(p) \leq 5\rho(x)$$

Less demanding oracle

$\mathrm{Vor}^{\pm}_{|S}(\mathcal{P})$ = edges of $\mathrm{Vor}(\mathcal{P})$ that intersect S an odd number of times

if $S = f^{-1}(x)$, deciding whether an edge $e = [pq]$ belongs to $\mathrm{Vor}^{\pm}_{|S}(\mathcal{P})$ reduces to evaluating the sign of f at p and q

The isotopy proof still holds

Computing $\mathrm{lfs}(x)$ is difficult

Computing $\mathrm{rch}(S) = \inf_{x \in S} \mathrm{lfs}(x)$ is much easier
$\mathrm{rch}(S)$ is either

- a local minimum of the smallest radius of curvature or
- the radius of a sphere with a diameter binormal to S

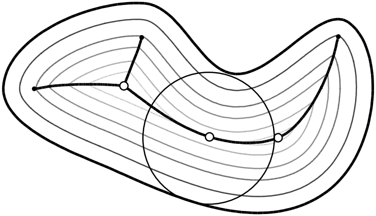

Computing $\operatorname{lfs}(x)$ is difficult

Computing $\operatorname{rch}(S) = \inf_{x \in S} \operatorname{lfs}(x)$ is much easier
$\operatorname{rch}(S)$ is either

- a local minimum of the smallest radius of curvature or
- the radius of a sphere with a diameter binormal to S

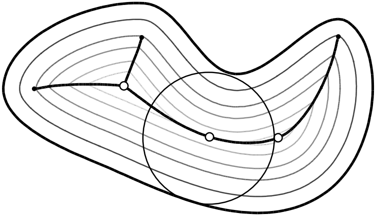

$$
\begin{aligned}
&F(p) = 0 \\
&F(q) = 0 \\
&(p-q) \times \nabla f(p) = 0 \\
&(p-q) \times \nabla f(q) = 0 \\
&\lambda\,(p-q)^2 = 1
\end{aligned}
$$

Applications

- Implicit surfaces $f(x, y, z) = 0$
- Isosurfaces in a 3d image (Medical images)
- Triangulated surfaces (Remeshing)
- Point sets (Surface reconstruction)

Results on smooth implicit surfaces

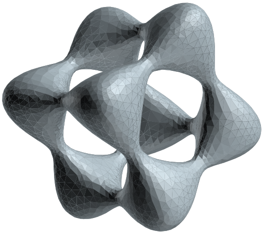

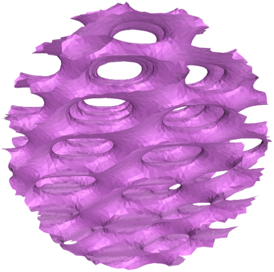

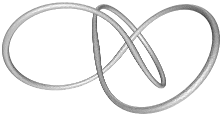

Pentium IV, 3.6GHz

Surface	Output size	Combinatorial	Bipolar oracle	time	CPU time
Tangle cube	4,242	8.31%	0.81%	8.52%	2.42 s
Trefoil	8,317	12.54%	0.93%	13.47%	5.14 s

% are wrt Del($\mathcal{P}$)

Comparison with the Marching Cube algorithm

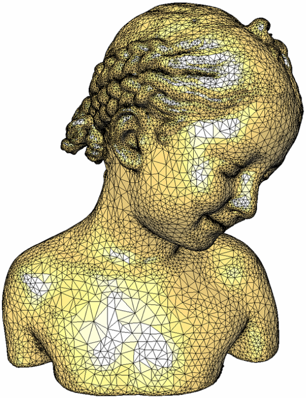

Delaunay Refinement

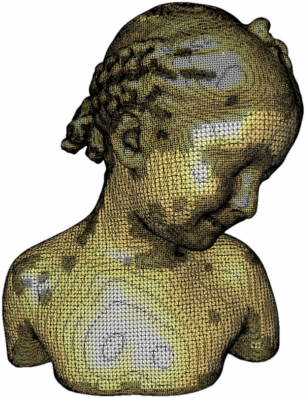

Marching Cube

Comparison with the Marching Cube algorithm

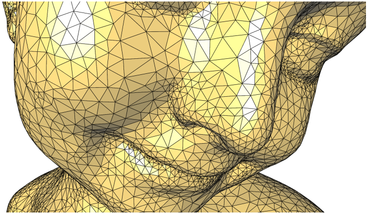

Delaunay refinement

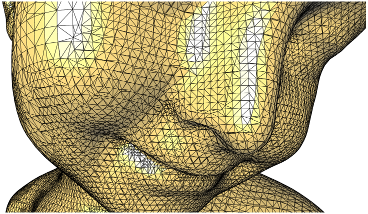

Marching cube

Contouring isosurfaces in 3D images

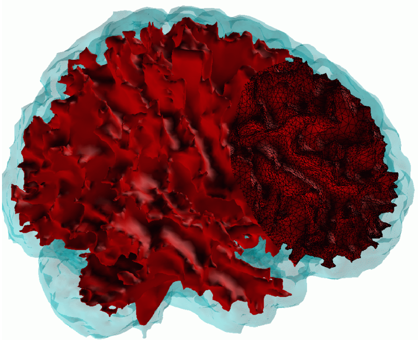

Collaboration with Asclepios, Caïman and Odyssée INRIA project-teams

Non smooth surfaces

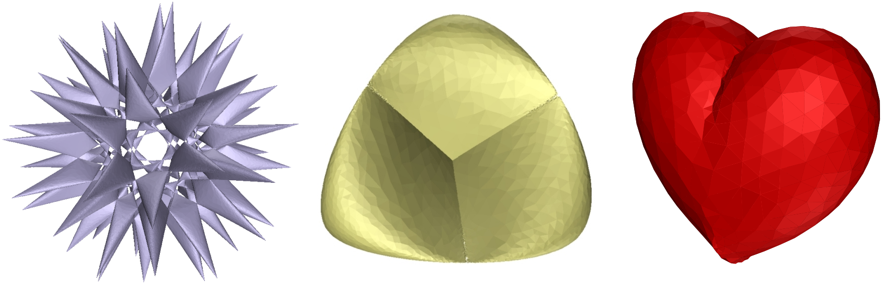

Extensions to

k-Lipschitz surfaces	[B. Oudot 06]
piecewise smooth surfaces	[Dey et al., Rineau & Yvinec 07]

Remeshing Polyhedral Surfaces

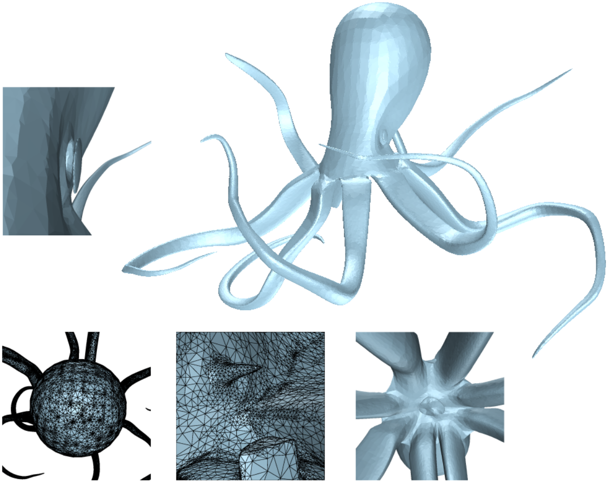

Point set surfaces

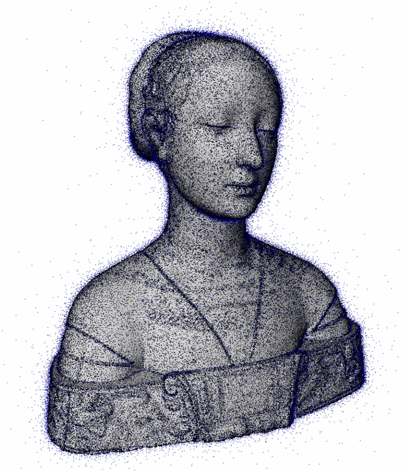

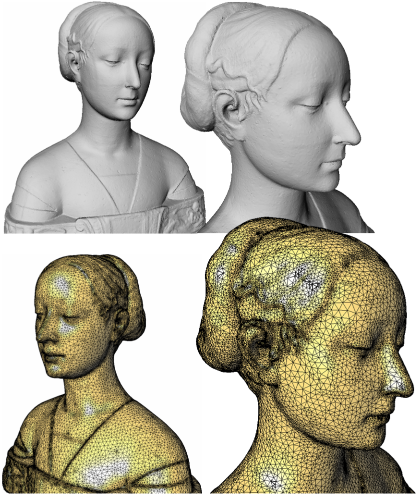

Courtesy of P. Alliez

Extension to non-binary datasets

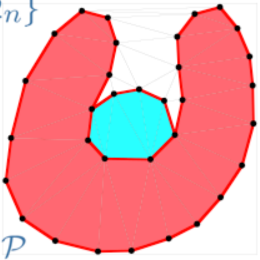

- Partition of space $\mathcal{P} = \{\Omega_0, \Omega_1, \ldots, \Omega_n\}$
  $\mathbb{R}^3 = \sqcup_{i \in \{0,\ldots,n\}} \Omega_i$
- Boundaries $\Gamma = \cup_i \delta\Omega_i$
- E a “good” point sample of Γ
- **Partition of Delaunay tetrahedra induced by** $\mathcal{P}$

$$\mathrm{Del}|_{\mathcal{P}}(E) = \{\mathrm{Del}|_{\Omega_0}(E), \ldots, \mathrm{Del}|_{\Omega_n}(E)\}$$

= each tetrahedron is labeled with the tissue its circumcenter belongs to

➔ **Consistent triangular and tetrahedral meshes**

- It “suffices” to generate E !

Results: Surface meshing with tissue-dependent resolution

- 72 tissues
- 112K vertices
- 228K b. facets
- 728K tets
- 340 seconds
- Criteria:
  - min angle>30°
  - cortex size<1mm
  - others size<2mm

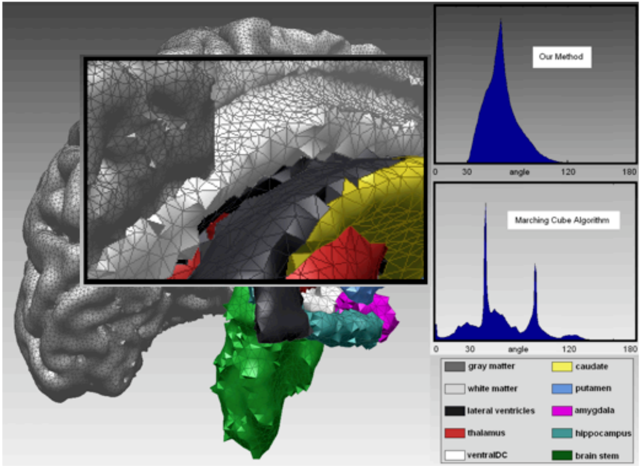

Results: Uniform surface meshing

- 15 tissues
- 11K vertices
- 23K b. facets
- 73K tets
- 35 seconds
- Criteria:
  - min angle>30°
  - size<1mm

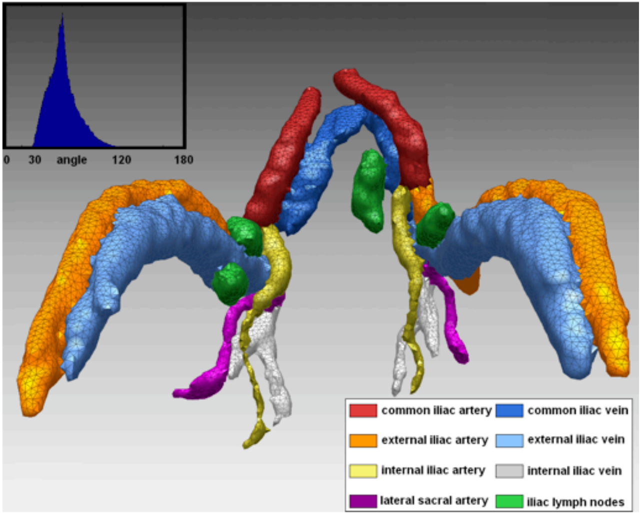

Multi-view 3D reconstruction

[Aganj et al 2007]

- **Multi-view 3D reconstruction:** Recovering the 3D shape of an object from several images taken from different viewpoints

- **Shape-from-silhouettes:** Only uses foreground/background segmentations

Problem statement

- **Visual hull:** Intersection of the visual cones of the different cameras
  - ☺ Very fast
  - ☹ Coarse reconstruction: OK for rendering

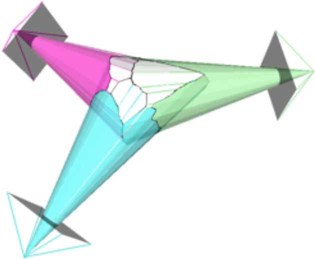

- Previous work
  - *Volumetric methods*
    - ☺ Fast, straightforward
    - ☹ High memory cost
    - ☹ Huge output mesh
    - ☹ Discretization artefacts (stair-casing effect)
  - *Exact polyhedral intersection* (Franco & Boyer, 2003)
    - ☺ Very fast
    - ☹ Low quality mesh
    - ☹ Noisy/inconsistent silhouettes ➡ numerical instabilities

Delaunay meshing approach

- **Multi-view oracle**:

  A Delaunay tetrahedron is part of the approx. visual hull
  $\Leftrightarrow$
  Its circumcenter projects inside the silhouette in all views

- **Refinement criterion = reprojection error**

  A boundary facet is good
  $\Leftrightarrow$
  All projections are closer from the silhouette boundary than a user-defined threshold in all views

- Pros and cons
  - ☹ Slower
  - ☺ High quality mesh
  - ☺ Reprojection error control (one-sided)

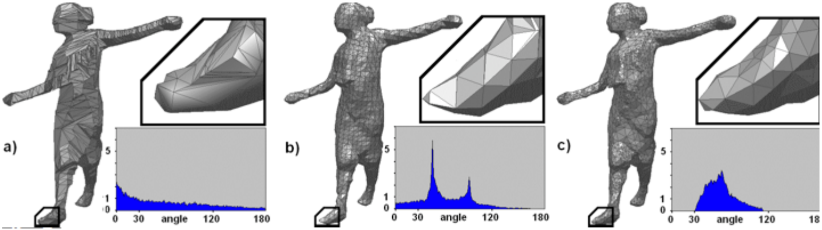

Polyhedral intersection Marching cubes Delaunay meshing

NB: same mesh size

Spatio-temporal visual hull using 4D Delaunay meshing

[Aganj et al 2007]

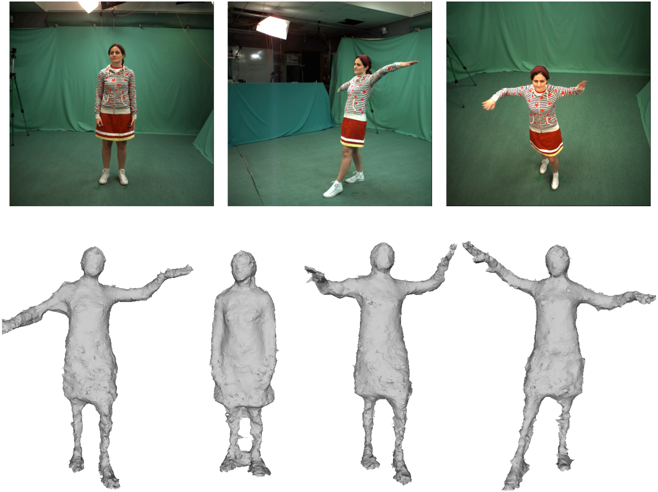

Approach

- Extend the previous algorithm to compute 4D visual hulls
  $\Rightarrow$ 4D representation of the scene

Advantages over frame-by-frame computations

- Exploits time redundancy
- Continuous representation, allowing spatio-temporal smoothing
- Reduction of flickering artefacts in synthesized views
- Handles naturally topological changes along time

– Requires an efficient implementation of DT in $\mathbb{R}^4$

Spatio-temporal multi-view oracle

A Delaunay pentatope is part of the approx. spatio-temporal visual hull	$\Leftrightarrow$	its circumcenter projects inside the time-interpolated silhouette in all views

Spatio-temporal multi-view oracle

A Delaunay pentatope is part of the approx. spatio-temporal visual hull $\Leftrightarrow$ its circumcenter projects inside the time-interpolated silhouette in all views

Refinement criterion = reprojection error

Accounts for spatio-temporal curvature
e.g. uniform motion $\rightarrow$ coarser resolution

Spatio-temporal multi-view oracle

A Delaunay pentatope is part of the approx. spatio-temporal visual hull $\Leftrightarrow$ its circumcenter projects inside the time-interpolated silhouette in all views

Refinement criterion = reprojection error

Accounts for spatio-temporal curvature
e.g. uniform motion $\rightarrow$ coarser resolution

Computing 3D temporal slices

- The boundary B of STVH is a set of tetrahedra $\subset \mathbb{R}^4$
- March on the tet of B to compute the intersection of B with a hyperplane $t = \text{constant}$
- A 3-facet of $B \rightarrow \emptyset$, a triangle or a quad

Results on real data

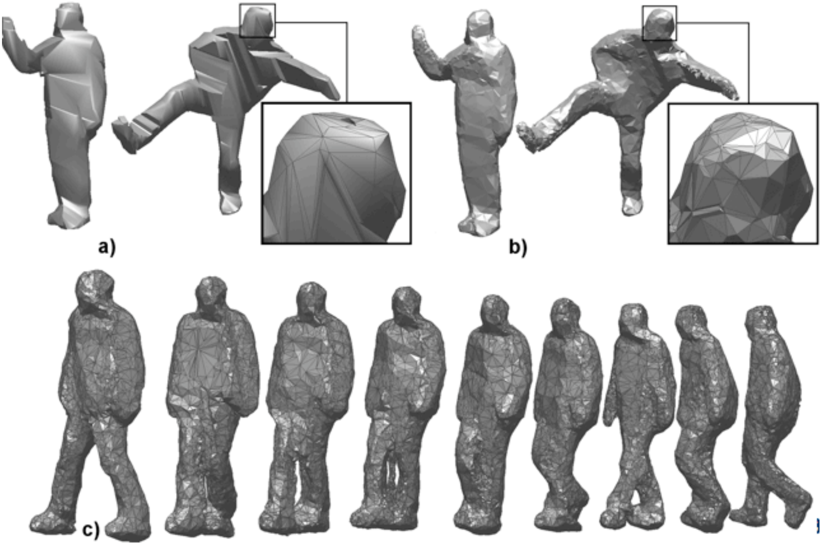

Delaunay deformable models [Pons & B. 2007]

Problem

To model moving surfaces undergoing large deformations and topology changes

Our approach

- Represent the interface by a triangular mesh embedded in the restricted 3D DT of interface points
- Update the mesh at each time step by updating the RDT

The algorithm at time step n

input: $\mathcal{P}_n = \textit{sample}$, Ω_n, $D_n = \mathrm{Del}_{|\Omega_n}(\mathcal{P}_n)$

1. Move the points of $\mathcal{P}_n$
2. Adapt the sample resolution $\rightarrow \mathcal{P}_{n+1}$
   2.1 replace a too short edge of D_n by its midpoint
   2.2 add the midpoint of a too long edge
3. Deal with topology changes
   3.1 discard tetrahedra that have been inverted
   $\Omega_{n+1} = \bigcup\{\tilde{\tau} | \tau \in D_n, \texttt{orient}(\tilde{\tau}) \times \texttt{orient}(\tau) < 0\}$
   3.2 $D_{n+1} = \mathrm{Del}_{|\Omega_{n+1}}(\mathcal{P}_{n+1})$
   3.3 remove the vertices whose incident simplices all share the same label

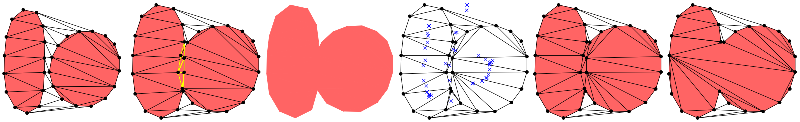

(a) (b) (c) (d) (e) (f)

A sphere deforming into a torus

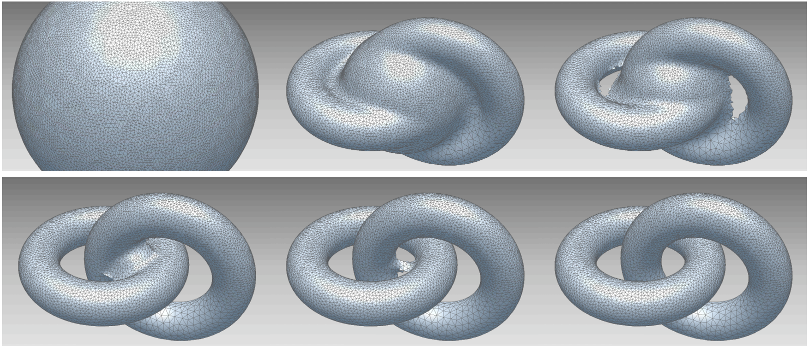

Application to image segmentation

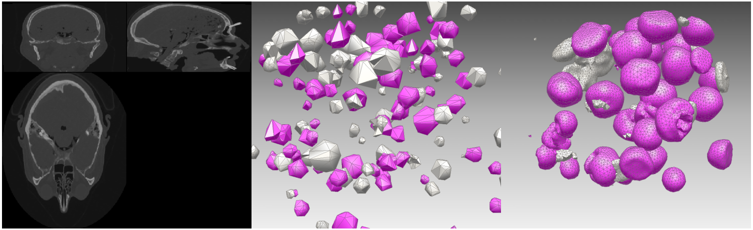

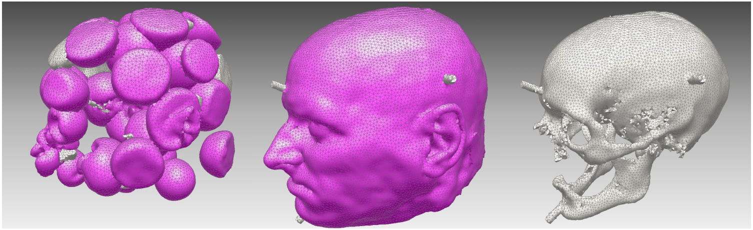

Collaboration with J-P. Pons (CERTIS)

Meshing 3D domains by Delaunay refinement

The surface mesher

- inserts points on the surface S
- triangulates the ambient space and extracts the Delaunay triangulation restricted to S
- controls the shape of the triangles of $\mathrm{Del}_{|S}(\mathcal{P})$

Meshing 3D domains by Delaunay refinement

The surface mesher

- inserts points on the surface S
- triangulates the ambient space and extracts the Delaunay triangulation restricted to S
- controls the shape of the triangles of $\mathrm{Del}_{|S}(\mathcal{P})$

Hence

- it can triangulate the domain O bounded by S at no additional cost
- but does not provide control on the shape of the tetrahedra inside O

Meshing volumes with curved boundaries [Oudot et al. 05]

3-d mesh refinement algorithm

1. Run the surface meshing algorithm

2. Insert points inside O to remove the bad elements of $\mathrm{Del}_{|S}(\mathcal{P})$ and $\mathrm{Del}_{O}(\mathcal{P})$

Sizing field

$\psi(x)$ defined over O

Basic procedures

`refine_face`(f) : insert c_f, the center of the surface Delaunay ball circumscribing f

`refine_tet`(t) : insert c_t, the center of the ball circumscribing t

Sizing field

$\psi(x)$ defined over O

Basic procedures

`refine_face`(f) : insert c_f, the center of the surface Delaunay ball circumscribing f

`refine_tet`(t) : insert c_t, the center of the ball circumscribing t

Bad elements

bad facet f : $r_f > \alpha\psi(c_f)$ or f has a vertex $\notin S$

bad tet. t : tetrahedron whose circumscribing ball has radius $r_t > \psi(c_t)$ or a radius-edge ratio $> \rho$

Step 2 : Apply the following rules in order

1. if $f \in \mathrm{Del}_{|S}(\mathcal{P})$ is a bad facet, `refine_face`(f)
2. if $t \in \mathrm{Del}_{O}(\mathcal{P})$ is a bad tetrahedron,
   2.1 if c_t is included in a surface Delaunay ball B_f, `refine_face`(f)
   2.2 else `refine_tet`(t)

Step 2 : Apply the following rules in order

1. if $f \in \mathrm{Del}_{|S}(\mathcal{P})$ is a bad facet, `refine_face`(f)
2. if $t \in \mathrm{Del}_O(\mathcal{P})$ is a bad tetrahedron,
   2.1 if c_t is included in a surface Delaunay ball B_f, `refine_face`(f)
   2.2 else `refine_tet`(t)

Properties

1. For appropriate α and ρ, the algorithm terminates
2. $\mathrm{Del}_{|S}(\mathcal{P}) = \mathrm{Del}_{|S}(\mathcal{P} \cap S)$ (cf. def of bad facet)
3. hence $\mathrm{Del}_S(\mathcal{P})$ is a 2-triangulation isotopic to S
   $\mathrm{Del}_O(\mathcal{P})$ is a 3-triangulation isotopic to O

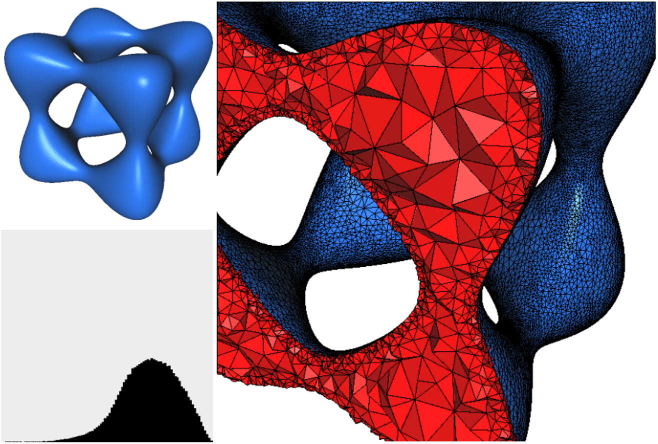

Non uniform mesh

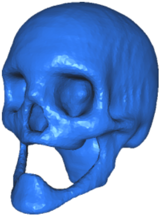
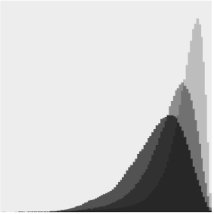
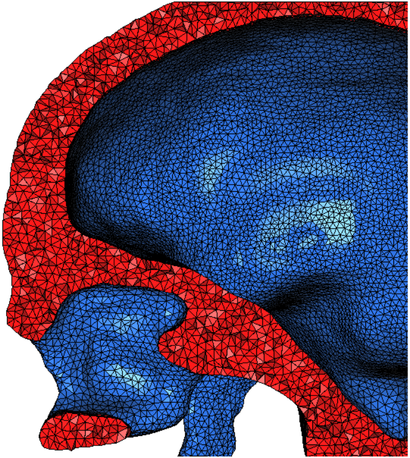

Uniform mesh

33,012 initial vertices, 2,471 + 53,762 new vertices
20s (Pentium IV, 1.7 GHz)

Multibody mesh generation from segmented images

input : a segmented 3D image (each voxel has a label)

label of a tetrahedron: label of its circumcenter

boundary facet : dual to a Voronoi edge whose endpoints have two different labels

Multibody mesh generation from segmented images

input : a segmented 3D image (each voxel has a label)

label of a tetrahedron: label of its circumcenter

boundary facet : dual to a Voronoi edge whose endpoints have two different labels

- We mesh simultaneously the various tissues using Delaunay refinement
- The boundary facets produce a good approximation of the interfaces

  all boundary surfaces are water tight
  and don't intersect each other

- The tetrahedra of a given label produce a good approximation of the associated tissue

Multibody mesh generation from segmented images

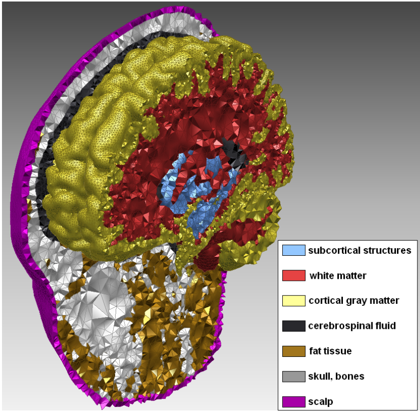

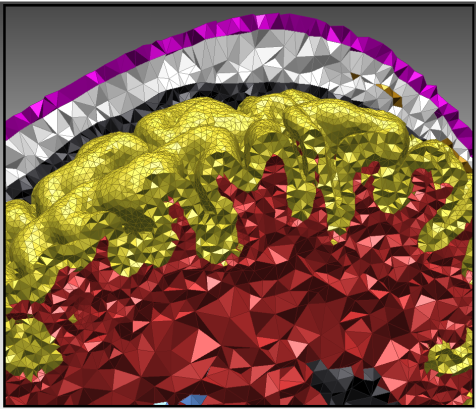

Tissue-dependent resolution

77 tissues, 389K vertices, 536K boundary facets, 536K tets
23 mn

Collaboration with CERTIS

Zigzaging effect along sharp features

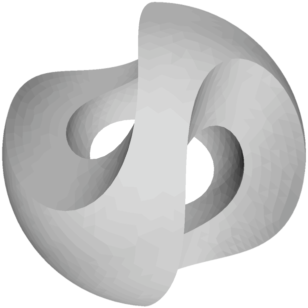

Input domain

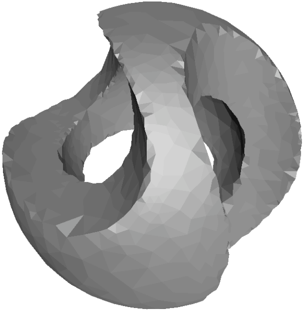

Output mesh

Zigzaging effect along 1-junctions between 3 or more tissues

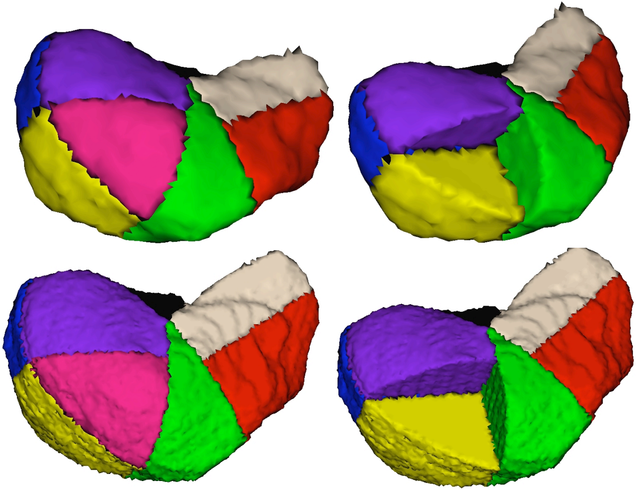

2 792 vertices,
5 681 triangles

20mm edge length

36 208 vertices
74 287 triangles

10mm edge length

Meshing 3D domains with piecewise smooth boundaries [Dey & Levine]

Protecting balls

- centered on the sharp features F of S
- B cannot contain the center of $B' \neq B$
- the balls cover F
- no 3 balls intersect

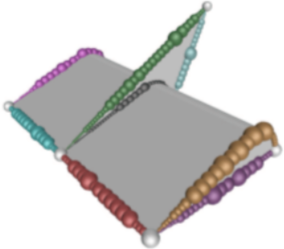

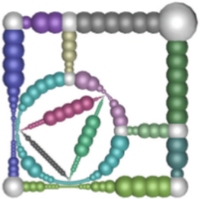

Meshing 3D domains with piecewise smooth boundaries [Dey & Levine]

Protecting balls

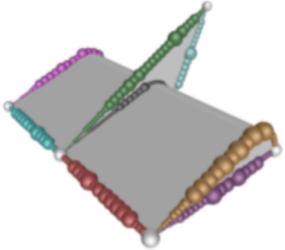

- centered on the sharp features F of S
- B cannot contain the center of $B' \neq B$
- the balls cover F
- no 3 balls intersect

Algorithm

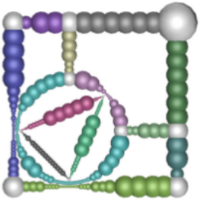

- Use the weighted DT
- Insert the protecting balls first
- Insert unweighted points inside O as usual

Meshing 3D domains with piecewise smooth boundaries

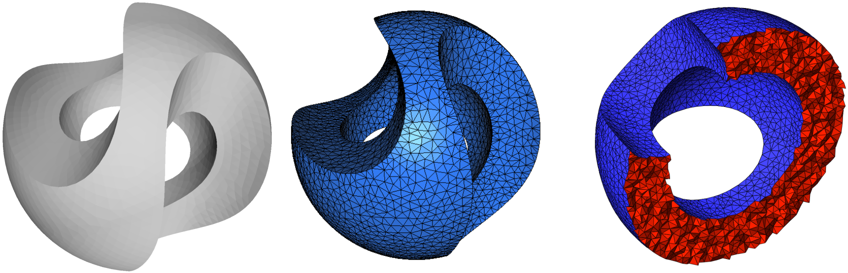

6 052	vertices
37 106	cells
$8,87^o$	smallest dihedral angle

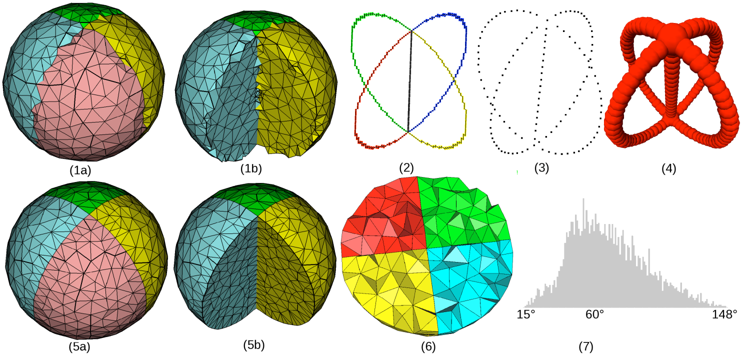

Figure 2: *(1) Delaunay refinement 3D mesh. (2) Multi-material junctions: five 1-junctions and two 0-junctions. (3) Sampled points on junctions. (4) Protecting balls. (5) Feature preserving Delaunay refinement 3D mesh. (6) A cut of the tetrahedral mesh. (7) Histogram of the dihedral angles.*

[Boltcheva et al. 2009]

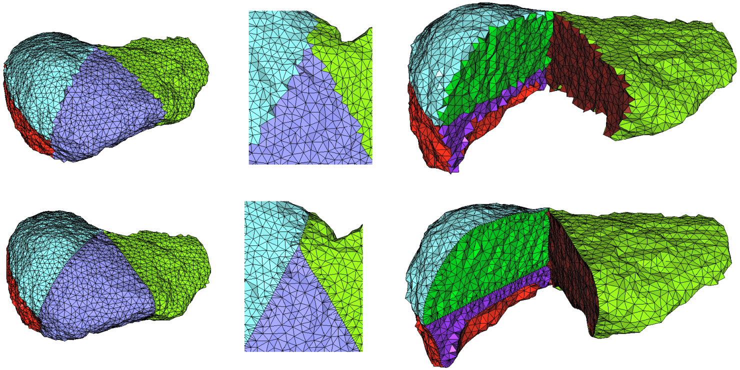

Figure 7: *Meshes generated from a segmented liver image representing 4 anatomical liver regions. The first row shows meshes obtained with the usual Delaunay refinement algorithm. The second row shows meshes generated with our feature preserving extension. The 3rd column shows some internal interfaces between the anatomical regions.*

[Boltcheva et al. 2009]

CGAL mesh generator www.cgal.org

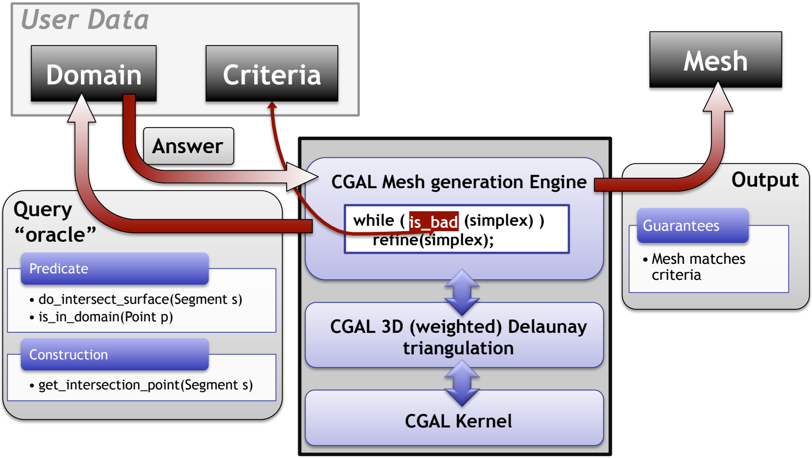

- core mesh generation algorithm independent from the input domain representation

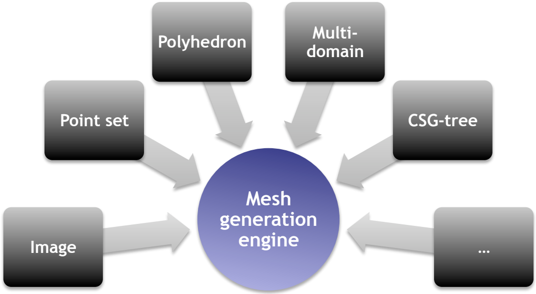

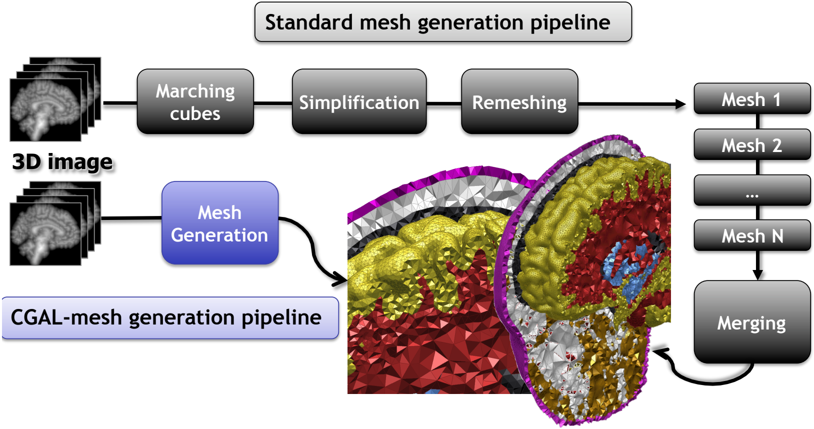
Standard mesh generation pipeline
Marching cubes
Simplification
Remeshing
Mesh 1
Mesh 2
...
Mesh N
Merging
3D image
Mesh Generation
CGAL-mesh generation pipeline

Acknowledgments

Students and colleagues

S. Oudot	Surface mesh generation [2005]
L. Rineau	3D domains mesh generation [2007]
C. Wormser	Anisotropic mesh generation [2008]
J. Tournois	Mesh optimization [2009]
J-P. Pons	Applications in computer vision
D. Boltcheva	Applications in medical imaging
N. Salman	Applications to surface reconstruction
P. Alliez, L. Rineau, S. Tayeb, M. Yvinec	CGAL-mesh

Acknowledgments

Projects

The CGAL library : http://www.cgal.org

IST Programme of the EU :

- ECG Project (Effective Computational Geometry for Curves and Surfaces) http://www-sop.inria.fr/prisme/ecg/
- ACS Project (Algorithms for complex shapes) http://acs.cs.rug.nl/

ADT CGAL-mesh