Smooth manifold reconstruction from point clouds

Jean-Daniel Boissonnat
INRIA Sophia-Antipolis

Winter School on Algorithmic Geometry
ENS-Lyon
January 2010

From unorganized 3D point clouds to triangulated surfaces : how to connect the dots ?

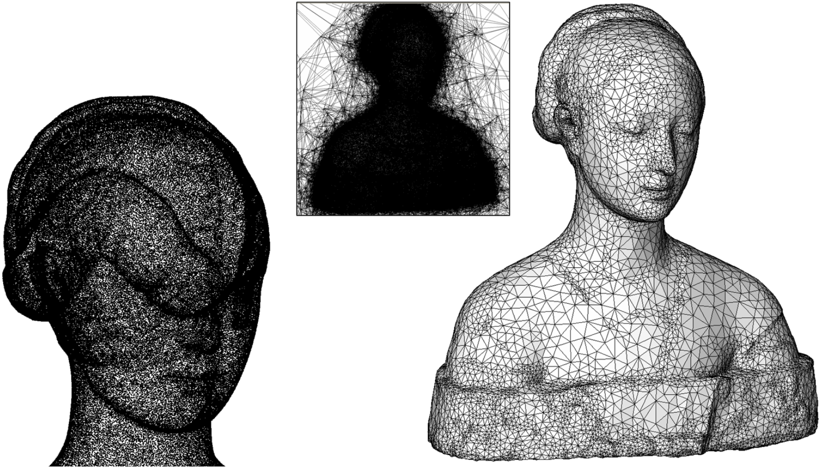

Overview

- Preliminaries : normal estimation and poles
- Combinatorial methods
- Implicit methods
- Continuation methods
- Multi-scale approaches

Estimating normal directions

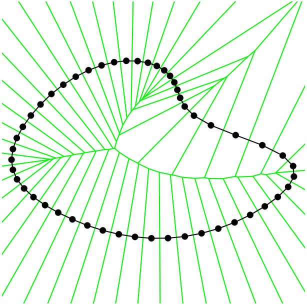

Voronoi regions are elongated in the normal direction

$V(p_i)$ contains the center m_{p_i} of the medial ball tangent to S at p_i

Pole of p_i : vertex v of $V(p_i)$ furthest from p_i

$$\|v - p_i\| \geq \|p_i - m_{p_i}\| \geq \mathrm{lfs}(x)$$

Error bound on the normal direction

$\mathcal{P} = \varepsilon$-sample

v pole of x : $\|v - x\| \geq \mathrm{lfs}(x)$

$\|i - x\| \leq \varepsilon \mathrm{lfs}(i) \leq \frac{\varepsilon}{1-\varepsilon} \mathrm{lfs}(x)$
(lfs 1-Lipschitz)

$$\angle(\vec{n}(x), \vec{vx}) = \alpha + \beta \\ \leq 2 \arcsin \frac{\varepsilon}{1-\varepsilon} \\ \approx 2\varepsilon$$

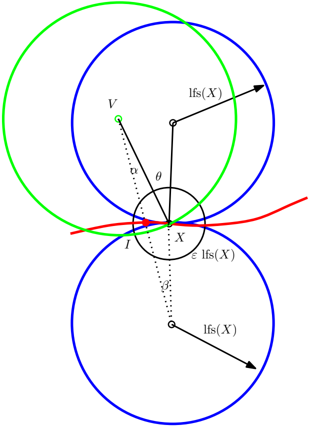

Orientation of the normals

orientation
may be given by the
scanning device

otherwise :

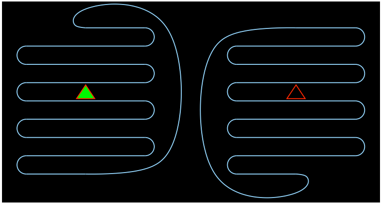

- orient the normals at the vertices of $\mathrm{conv}(\mathcal{P})$ (easy)
- then propagate the labels coherently
  - p_i close to $p_j \Rightarrow n_{p_i} \cdot n_{p_j} > 0$
  - heuristics : walk along the MST of $\mathcal{P}$ [Hoppe 92]

Warning !

The problem is global : local methods are non robust in the presence of

Slivers

- Crust and power crust [Amenta et al.]

Noise, sharp features and undersampling

- Spectral surface reconstruction [Kolluri et al. 2006]
- Voronoi PCA [Alliez et al. 2007]

[Mérigot 2009]

Power crust algorithm

[Amenta et al.]

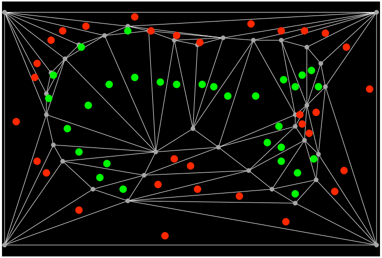

1. For each $p_i \in \mathcal{P}$, compute
   - its pole m_{p_i} and $w_i = \|m_{p_i} - p_i\|$
   - the secondary pole :
     $m'_{p_i} = \arg\min_{v \in V(p_i)} \left((m_{p_i} - p_i) \cdot (m'_{p_i} - p_i)\right)$
     and $w'_i = \|m'_{p_i} - p_i\|$
2. Compute the weighted Voronoi (power) diagram $WVor(WP)$ where $WP = \bigcup_{i=1..n}\{(m_{p_i}, w_i)\} \bigcup \{(m'_{p_i}, w'_i)\}$
3. Label the elements of WP as inside or outside
4. Output the facets of $WVor(WP)$ incident to cells with opposite labels

Power crust algorithm

[Amenta et al.]

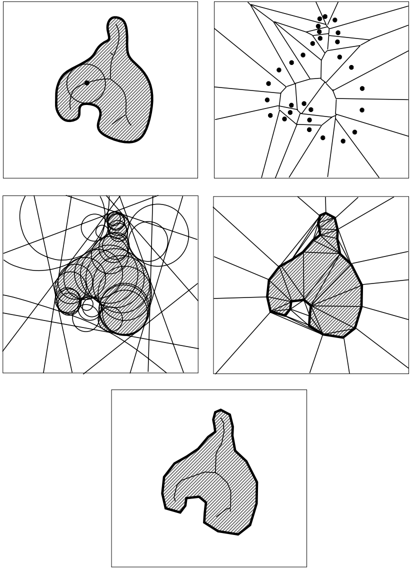

Global labelling of poles using a spectral approach

Kolluri, Shewchuk, O'Brien

Spectral surface reconstruction from noisy point clouds

Siggraph 2004

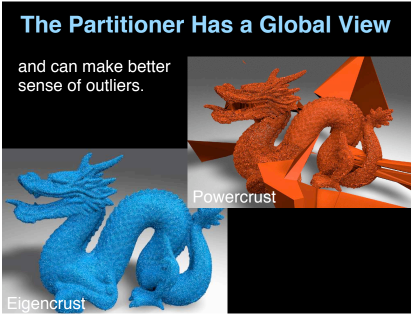

[Kolluri, Shewchuk, O'Brien]

Surface reconstruction : the implicit approach

1. Compute a function over $\mathbb{R}^3$ whose zero-set Z either interpolates or approximates E
2. mesh Z

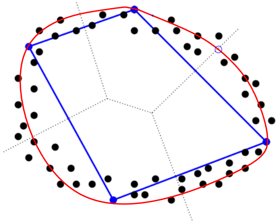

Various implicit functions

- interpolation of the sign distance functions to the tangent planes at the sample points [Hoppe 92], [B. & Cazals 00]
- moving least squares (MLS) [Levin 03]
- radial basis functions [Carr et al. 01]
- Poisson based approach [Khazdan et al. 06]
- spectral method [Alliez et al. 07]

Interpolating distance functions

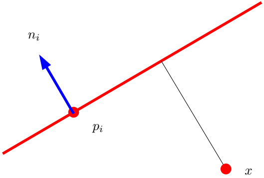

$$h_i(x) = (x - p_i) \cdot n_i$$

assuming we know or can estimate the normals n_i (direction and orientation) to S at p_i

Interpolating distance functions

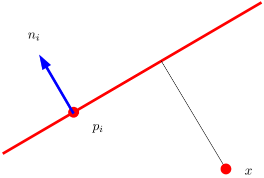

$$h_i(x) = (x - p_i) \cdot n_i$$

assuming we know or can estimate the normals n_i (direction and orientation) to S at p_i

Functional interpolant $\quad h(x) = \sum_i \lambda_i(x) h_i(x)$

where $\forall x,\ \sum_{i=1}^{n} \lambda_i(x) = 1$ (partition of unity)

Interpolating distance functions

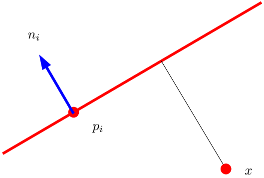

$$h_i(x) = (x - p_i) \cdot n_i$$

assuming we know or can estimate the normals n_i (direction and orientation) to S at p_i

Functional interpolant $\quad h(x) = \sum_i \lambda_i(x) h_i(x)$

where $\forall x, \ \sum_{i=1}^{n} \lambda_i(x) = 1$ (partition of unity)

Implicit surface $h^{-1}(0)$

Interpolating distance functions

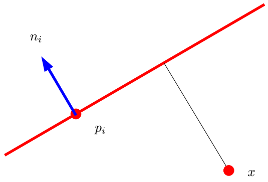

$$h_i(x) = (x - p_i) \cdot n_i$$

assuming we know or can estimate the normals n_i (direction and orientation) to S at p_i

Functional interpolant $h(x) = \sum_i \lambda_i(x) h_i(x)$

where $\forall x,\ \sum_{i=1}^{n} \lambda_i(x) = 1$ (partition of unity)

Implicit surface $h^{-1}(0)$

Properties

- $\lambda_i(p_j) = \delta_{ij} \ \Rightarrow \ h(p_i) = h_i(p_i) = 0$
- if the λ_i are C^k continuous, $h^{-1}(c)$ is a C^k surface for almost all c (by Sard's th.)

Choice of $\lambda_i(x)$

Some examples

- Nearest neighbor (NN) : $\lambda_i(x) = 1$, $\lambda_{j\neq i}(x) = 0$ if $x \in V(p_i)$
  Not continuous, support $= V(p_i)$
- Radial basis function : $\lambda_i(x) = \exp^{-\beta\|x-p_i\|^2}$, $\beta > 0$
  C^∞, non compact support
- Natural neighbor coordinates
  C^1-continuous, compact support

Natural Neighbor Coordinates

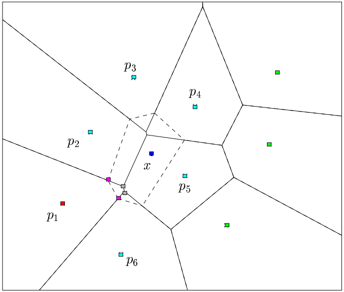

[Sibson 80]

Natural neighbors

For $x \in \mathrm{conv}(E)$

$\mathrm{Nat}(x) = \{p_i : w_i(x) \neq 0\}$

Natural Neighbor Coordinates

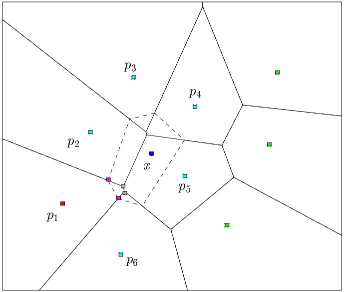

[Sibson 80]

Natural neighbors

For $x \in \mathrm{conv}(E)$

$\mathrm{Nat}(x) = \{p_i : w_i(x) \neq 0\}$

Laplace's coordinates

$$\lambda_i(x) = \frac{v_i(x)}{\sum_i v_i(x)} \text{ with } v_i(x) = \frac{\mathrm{vol}(V^+(x,p_i))}{\|x - p_i\|}, \qquad \lambda_i(p_i) = 1,\ \lambda_j(p_i) = 0$$

Natural Neighbor Coordinates

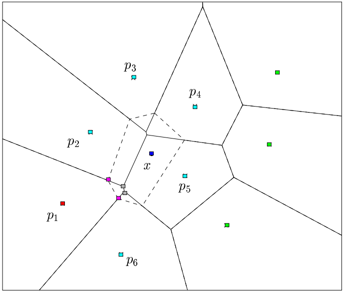

[Sibson 80]

Natural neighbors

For $x \in \mathrm{conv}(E)$

$\mathrm{Nat}(x) = \{p_i : w_i(x) \neq 0\}$

Laplace's coordinates

$$\lambda_i(x) = \frac{v_i(x)}{\sum_i v_i(x)} \text{ with } v_i(x) = \frac{\mathrm{vol}(V^+(x,p_i))}{\|x - p_i\|}, \qquad \lambda_i(p_i) = 1,\ \lambda_j(p_i) = 0$$

Sibson's coordinates

$$\sigma_i(x) = \frac{w_i(x)}{\sum_i w_i(x)} \text{ with } w_i(x) = \mathrm{vol}(V^+(p_i) \cap V^+(x))$$

Equivalently $\mathrm{Nat}(x) = \{$ vertices of star(x) in $Del^{+}(E)$ $\}$
$= \{$ vertices of the Delaunay simplexes whose circumscribing ball $\ni x\}$

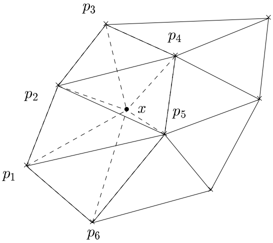

Equivalently $\mathrm{Nat}(x) = \{$ vertices of star(x) in $Del^+(E)$ $\}$
$= \{$ vertices of the Delaunay simplexes whose circumscribing ball $\ni x\}$

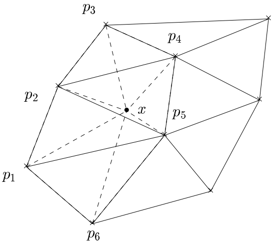

$\mathbb{R}^d/\mathrm{Nat}$ is the arrangement of the Delaunay spheres

Equivalently $\mathrm{Nat}(x) = \{$ vertices of star(x) in $Del^+(E)$ $\}$
$= \{$ vertices of the Delaunay simplexes whose circumscribing ball $\ni x\}$

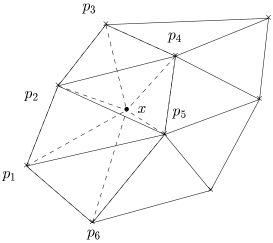

$\mathbb{R}^d/\mathrm{Nat}$ is the arrangement of the Delaunay spheres

Support of λ_i or σ_i = union of the Delaunay balls incident to p_i

Properties of the Laplace coordinates $\lambda_i(x)$

$$\lambda_i(x) = \frac{v_i(x)}{\sum_i v_i(x)} \text{ with } v_i(x) = \frac{\mathrm{vol}(V^+(x,p_i))}{\|x-p_i\|}$$

1. Trivially $\sum_i \lambda_i(x) = 1$

2. The λ_i are continuous

3. when $x \to p_i$, $\quad v_i(x) \to \infty$
$\quad v_j(x) \to \mathrm{cst} < \infty \qquad (j \neq i)$

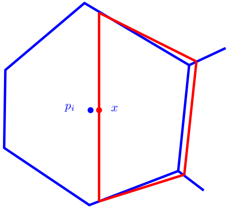

$$\Rightarrow \qquad \begin{array}{l} \lambda_i(x)) \to 1 \\ \lambda_j(x)) \to 0 \end{array}$$

4. $x = \sum_i \lambda_i(x)\, p_i$?

$$x = \sum_i \lambda_i(x)\, p_i$$

$x = \sum_i \lambda_i(x)\; p_i$

Minkowski's th.

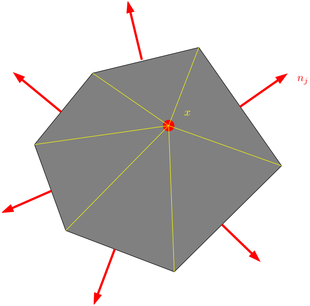

$$d\,\mathrm{vol}(P) = \sum_{j\in J} \mathrm{vol}(f_j)\;(x - p_j)\cdot n_j$$

$$\nabla \mathrm{vol}(P) = 0 \;\; \Rightarrow \sum_{j\in J} \mathrm{vol}(f_j)\; n_j = 0$$

$x = \sum_i \lambda_i(x)\, p_i$

Minkowski's th.

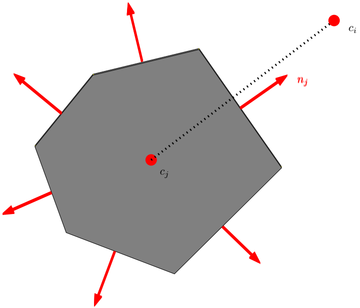

$$d\,\mathrm{vol}(P) = \sum_{j\in J} \mathrm{vol}(f_j)\,(x - p_j)\cdot n_j$$

$$\nabla \mathrm{vol}(P) = 0 \quad \Rightarrow \sum_{j\in J} \mathrm{vol}(f_j)\; n_j = 0$$

Applied to $V^+(x)$

Hence $\sum_i v_i \frac{x-p_i}{\|x-p_i\|} = 0 \iff \left(\sum_i \frac{v_i}{\|x-p_i\|}\right) x = \sum_i \frac{v_i}{\|x-p_i\|} p_i$

Properties of the Sibson's coordinates $\sigma_i(x)$

1. Trivially $\sum_i \sigma_i(x) = 1$

2. The σ_i are continuous

3. when $x \rightarrow p_i$, $\quad w_i(x) \rightarrow C \neq 0$
   $\quad w_j(x) \rightarrow 0 \quad (j \neq i)$

   $\Rightarrow \quad \sigma_i(x)) \rightarrow 1$
   $\quad \sigma_j(x)) \rightarrow 0$

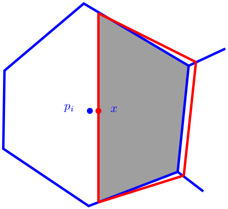

4. $x = \sum_i \sigma_i(x)\ p_i$

   same proof as for the Laplace coord. but in the space of spheres

$x = \sum_i \sigma_i(x)\, p_i$

Minkowski's th. applied to $\mathcal{V}_p = h_x^- \cap \{h_i^+\}$

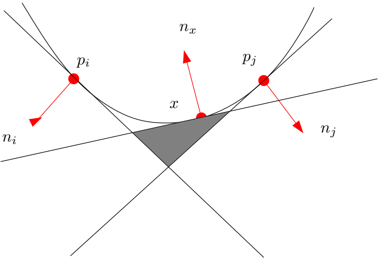

$x = \sum_i \sigma_i(x)\, p_i$

Minkowski's th. applied to $\mathcal{V}_p = h_x^- \cap \{h_i^+\}$

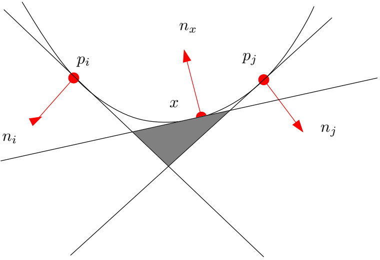

$$n_i = (2\,p_i, -1) \qquad \perp f_i$$
$$n_p = (-2\,x, 1) \qquad \perp f_x$$

$$\sum_i \mathrm{vol}(f_i)\ \frac{n_i}{\|n_i\|} + \mathrm{vol}(f_x)\ \frac{n_x}{\|n_x\|} = 0$$

$$w_i(x) = \mathrm{vol}(\,V^+(p) \cap V(p_i)) =$$
$$\mathrm{vol}(f_i)\ \frac{n_i}{\|n_i\|} \cdot (-i_{d+1}) = \frac{\mathrm{vol}(f_i)}{\|n_i\|}$$
$$w(x) = \mathrm{vol}(\,V^+(x)) = \frac{\mathrm{vol}(f_x)}{\|n_x\|}$$

$x = \sum_i \sigma_i(x)\, p_i$

Minkowski's th. applied to $\mathcal{V}_p = h_x^- \cap \{h_i^+\}$

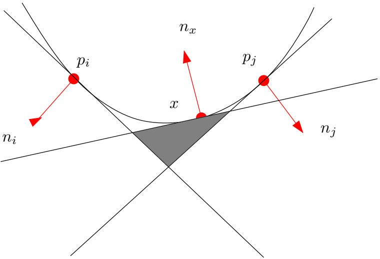

$$n_i = (2\,p_i, -1) \qquad \perp f_i$$
$$n_p = (-2\,x, 1) \qquad \perp f_x$$

$$\sum_i \text{vol}(f_i)\ \frac{n_i}{\|n_i\|} + \text{vol}(f_x)\ \frac{n_x}{\|n_x\|} = 0$$

$$w_i(x) = \text{vol}(\,V^+(p) \cap V(p_i)) =$$
$$\text{vol}(f_i)\ \frac{n_i}{\|n_i\|} \cdot (-i_{d+1}) = \frac{\text{vol}(f_i)}{\|n_i\|}$$
$$w(x) = \text{vol}(\,V^+(x)) = \frac{\text{vol}(f_x)}{\|n_x\|}$$

$$\sum_i w_i(x) n_i + w(x) n_x = 0 \quad \Longrightarrow \quad \sum_i w_i(x) p_i \;-\; w(x) x = 0$$

Strength of the implicit approach : Refining the mesh allows to remove singularities

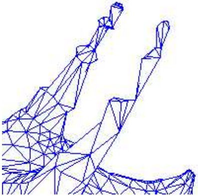

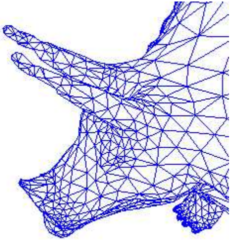

(Using Sibson's interpolation)

Certified algorithms and codes

- Cocone algorithm [Dey & al.]
- Geomagic Studio
- Natural neighbor interpolation of distance functions (Catia)
- Numerical approaches [Kolluri & al., Khazdan & al., Alliez & al.]
- CGAL-3.5 :
  Geometry Processing & Surface Reconstruction [Alliez & al.]

Smooth submanifold reconstruction

- Motivation and difficulties
- Continuation approach and intrinsic complexity
- Dimension estimation and multi-scale reconstruction

Motivations

- Theoretical foundations and guarantees
  (Sampling theory for geometric objects)
- Complexity issues and practical efficiency
- Higher dimensions
  - Reconstruction in 3D+t space
  - 6D phase space
  - Configuration spaces of robots, molecules...
  - Data analysis

The curses of Delaunay triangulations in higher dimensions

- Subdividing the ambient space is very costly
  - ? Can we store $\mathrm{Del}(\mathcal{P})$ in a more compact way?
  - ? Can we construct more local data structures?

- The restricted Delaunay triangulation is no longer a good approximation of the manifold even under strong sampling conditions
  - ? What else ?

Is Delaunay triangulation of practical use

in high dimensional spaces?

Combinatorial complexity of the Delaunay triangulation

- Worst-case: $O(n^{\lceil \frac{d}{2} \rceil})$
- Uniform distribution in $\mathbb{R}^d$: $c^d O(n)$ [Dwyer 1991]

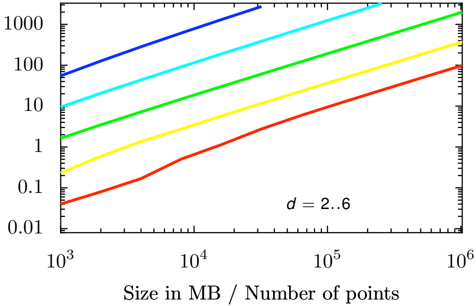

Algebraic issues

Deciding whether $x \in B(p_0, ..., p_d)$ reduces to evaluating the sign of

$$\sigma(p_0, ..., p_d, x) = \begin{vmatrix} 1 & ... & 1 & 1 \\ p_0 & ... & p_d & x \\ p_0^2 & ... & p_d^2 & x^2 \end{vmatrix}$$

σ is a polynomial of degree $d+2$ in the input variables
(its exact evaluation would require $(d+2)$-fold precision)

Implementation of the incremental algorithm [Hornus 2009]

- explicit storage of the *d*-simplices
- and of their adjacency graph

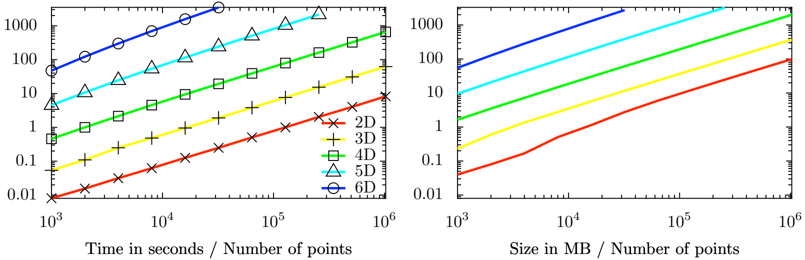

- Exact computing
- To come in CGAL 3.6

Compact representation

[B., Devillers, Hornus 2009]

Delaunay graph

- Store only the Delaunay graph (edges + vertices) + 1 d-simplex per vertex $\sigma(p)$
- `insert`(x) :
  1. locate the vertex of $\mathrm{Del}(\mathcal{P})$ nearest to x
  2. remove all d-simplices whose circ. ball $\ni x$ by walking from neighbor to neighbor in $\mathrm{Del}(\mathcal{P})$
  3. update the Delaunay graph by joining x to the vertices of the removed simplices

get_neighbor()

Input: the vertex-set of a Del-simplex σ and a vertex v *of* σ.
Output: the unique vertex w opposite to v.

Compute: $\displaystyle\bigcap_{p\in\sigma\setminus\{v\}} \mathrm{Neighbors}(p)$

Remove: vertices on the wrong side.

Use in_circle() predicate: to select the correct opposite vertex.

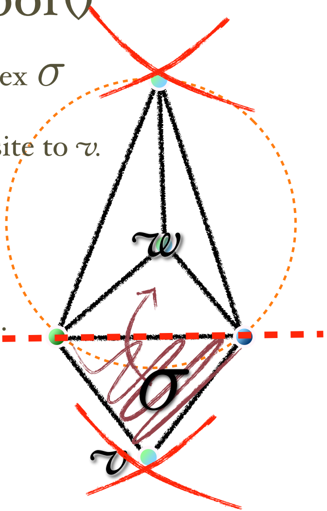

Performances of the Delaunay graph construction

1	Dimension	2	3	4	5	6
2	Number of input points	1024K	1024K	1024K	256K	32K
3	Size of the simplex-cache	1K	1K	10K	300K	1000K
4	Size of the conflict zone	4.1	21	134	940	6145
5	Calls to `neighbor(,)`	12.2	84.6	671.2	5631	43021
6	Number of candidates	2	2.6	4	6.7	11.6
7	Fast cache hit (non-*null* pointer)	56.6 %	57.5 %	54.6 %	55.5 %	54.3 %
8	Cache hit	37 %	39.6 %	40.1 %	42.3 %	43.1 %
9	Cache miss	6.4 %	2.9 %	5.3 %	2.2 %	2.6 %
10	Time ratio (`Del_graph/New_DT`)	6.1	5.7	6.0	6.5	8.1
11	Space ratio (`Del_graph/New_DT`)	2.7	1.7	0.6	0.2	0.1
12	Number of simplices per vertex	6	27 $(\times 4.5)$	157 $(\times 5.8)$	1043 $(\times 6.7)$	7111 $(\times 6.8)$
13	Number of edges per vertex	6	15.5 $(\times 2.6)$	36.5 $(\times 2.4)$	73 $(\times 2)$	164.6 $(\times 2.25)$

Current best: 100K 6D vertices: 105 millions simplices.
15 hours for the graph (approx 2:30 hours for full-D).

Space can be further improved using a compact representation of graphs

[Blandford et al 2003]

Is the restricted Delaunay triangulation a good approximation in high dimensional spaces?

Tangent space approximation

$\mathbb{M}$ is a smooth k-dimensional manifold embedded in $\mathbb{R}^d$

Bad news

The Delaunay triangulation restricted to $\mathbb{M}$ may be a bad approximation of the manifold even if the sample is dense

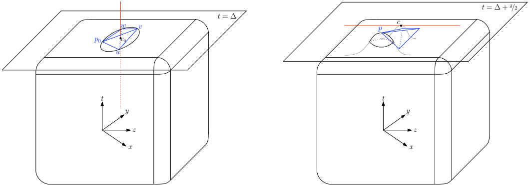

Tangent space approximation

$\mathbb{M}$ is a smooth k-dimensional manifold embedded in $\mathbb{R}^d$

Bad news

The Delaunay triangulation restricted to $\mathbb{M}$ may be a bad approximation of the manifold even if the sample is dense

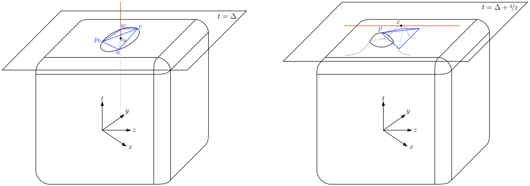

Good news

[Cheng et al. 2005]

If τ nor its faces are slivers, there exist a constant a_k (depending on σ_0) s. t. $\sin \angle(\mathrm{aff}(\tau), T_p) \leq a_k \varepsilon$

Slivers in higher dimensions

Sliver

Given constants ρ_0 and σ_0,
a j-simplex τ is called a sliver if

1. $j > 2$
2. $\rho_\tau = \frac{R_\tau}{L_\tau} \leq \rho_0$
3. $\sigma_\tau = \mathrm{vol}(\tau)/L_\tau^j < \sigma_0$
4. $\forall \sigma \subset \tau$, $\rho_\sigma \leq \rho_0$ and $\sigma_\sigma \geq \sigma_0$

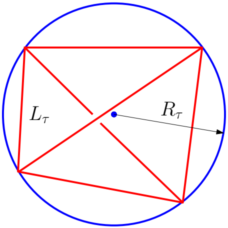

R_τ = radius of the circ. ball
L_τ = length of the shortest edge of τ

Slivers are fragile

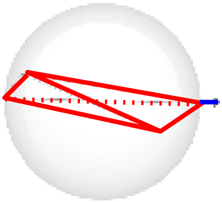

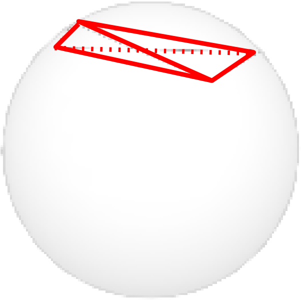

A small perturbation of one of its vertices may lead to

- a negative volume
- a big circumscribing sphere

vol(τ) is small

In both cases, the sliver is removed from DT

Two certified techniques to remove slivers

- Weight the vertices, replace the Delaunay triangulation by the regular triangulation of the WP [Cheng et al. 2000]
  - the points are not moved
  - must be used as a postprocessing step on the mesh
- Perturbe the vertices [Li 2001]
  - can be used during Delaunay refinement

Weighted Delaunay triangulation

Weight assignment

$$p_i \to \omega(p_i), \qquad \omega = (\omega(p_1), ..., \omega(p_n))$$

Weighted Voronoi diagram

$$V^\omega(p_i) = \{x : \|x - p_i\|^2 - \omega^2(p_i) \leq \|x - p_j\|^2 - \omega^2(p_j)\}$$

Weighted Delaunay triangulation

Weight assignment

$$p_i \to \omega(p_i), \qquad \omega = (\omega(p_1), \ldots, \omega(p_n))$$

Weighted Voronoi diagram

$$V^\omega(p_i) = \{x : \|x - p_i\|^2 - \omega^2(p_i) \leq \|x - p_j\|^2 - \omega^2(p_j)\}$$

Weighting a vertex

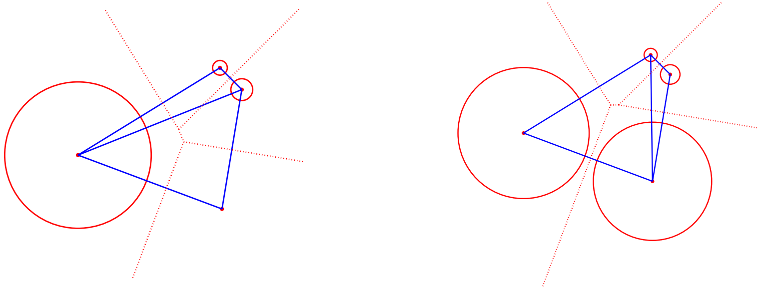

A continuation approach to manifold reconstruction

- Can we compute $\mathrm{Del}^{\omega}_{|S}(\mathcal{P})$ without computing $\mathrm{Del}(\mathcal{P})$?
- Can we avoid subdividing the embedding space and obtained an intrinsic complexity ?

Intrinsic complexity

Basic assumption in manifold learning

Data live in a low-dimensional manifold embedded in a high-dimensional space

Example 1 : human face images

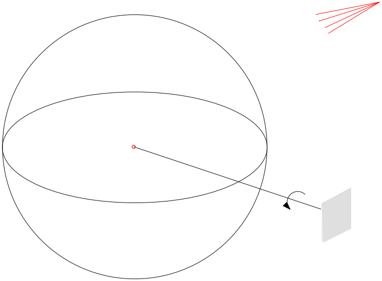

head = sphere
camera : 3 dof
light source : 2 dof

An image with N pixels $\rightarrow$ a point in $\mathbb{R}^N$

It is impossible to triangulate points in such a huge space !

Example 1 : human face images

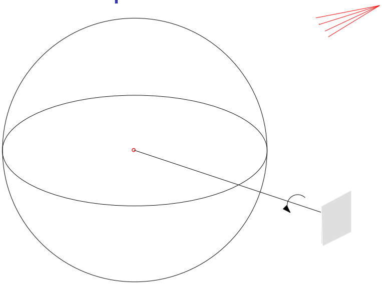

head = sphere
camera : 3 dof
light source : 2 dof

An image with N pixels $\rightarrow$ a point in $\mathbb{R}^N$

It is impossible to triangulate points in such a huge space !

Example 2 : points with unit normals

- $(p_i, \mathbf{n}_i) \in \mathcal{N} = \mathbb{R}^3 \times \mathbb{S}^2$
- The surface to be reconstructed is a 2-manifold embedded in $\mathcal{N}$

The tangential Delaunay complex

[Freedman 2002], [B.& Flottoto 2004], [B.& Ghosh 2009]

[Cheng, Dey, Ramos 2005]

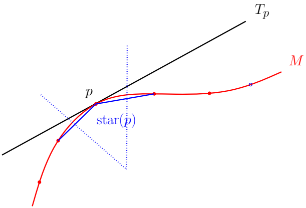

1. Construct the star of $p \in \mathcal{P}$ in the Delaunay triangulation $\mathrm{Del}_{T_p}(\mathcal{P})$ of $\mathcal{P}$ restricted to T_p
2. $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$: the set of stars of p, $p \in \mathcal{P}$

+ $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P}) \subset \mathrm{Del}(\mathcal{P})$
+ $\mathrm{star}(p)$, $\mathrm{Del}_{T_p}(\mathcal{P})$ and therefore $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$ can be computed without computing $\mathrm{Del}(\mathcal{P})$
– $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$ is not necessarily a triangulated manifold

Construction of $\mathrm{Del}_{T_p}(\mathcal{P})$

Given a k-flat H, $\mathrm{Vor}(\mathcal{P}) \cap H$ is a weighted Voronoi diagram

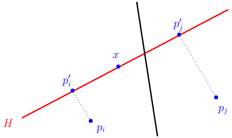

$$\|x - p_i\|^2 \leq \|x - p_j\|^2$$
$$\Leftrightarrow \quad \|x - p'_i\|^2 - \|p_i - p'_i\|^2 \leq \|x - p'_j\|^2 - \|p_j - p'_j\|^2$$

Construction of $\mathrm{Del}_{T_p}(\mathcal{P})$

Given a k-flat H, $\mathrm{Vor}(\mathcal{P}) \cap H$ is a weighted Voronoi diagram

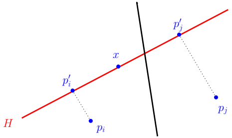

$$\|x-p_i\|^2 \leq \|x-p_j\|^2$$
$$\Leftrightarrow \quad \|x-p'_i\|^2 - \|p_i-p'_i\|^2 \leq \|x-p'_i\|^2 - \|p_j-p'_j\|^2$$

Corollary: construction of Del_{T_p}

1. project $\mathcal{P}$ onto T_p which requires $O(dn)$ time
2. construct $\mathrm{star}(p'_i)$ in $\mathrm{Del}^{\omega}(p'_i) \subset T_{p_i}$ where $\omega(p_i) = \|p_i - p'_i\|$
3. $\mathrm{star}(p_i) \overset{1-1}{\leftrightarrow} \mathrm{star}(p'_i)$

Inconsistent configurations $\phi = [p_1, ..., p_{k+2}]$

Definition

$\exists p_i, p_j, p_l \in \phi$ s.t.

1. $\tau = \phi \setminus \{p_l\} \in \mathrm{star}(p_i)$
   $\notin \mathrm{star}(p_j)$
2. τ nor its faces are slivers
3. $\mathrm{Vor}(p_l)$ is the first Voronoi cell intersected by $\vec{c_i c_j}$

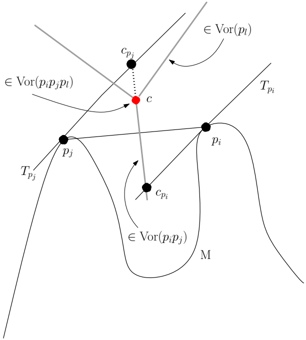

Inconsistent configurations $\phi = [p_1, ..., p_{k+2}]$

Definition

$\exists p_i, p_j, p_l \in \phi$ s.t.

1. $\tau = \phi \setminus \{p_l\} \in \text{star}(p_i)$
   $\notin \text{star}(p_j)$
2. τ nor its faces are slivers
3. $\text{Vor}(p_l)$ is the first Voronoi cell intersected by $\vec{c_i c_j}$

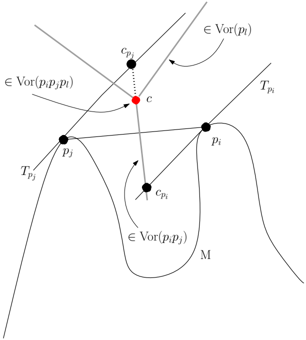

Observations

- c_i and c_j are close if p_i is close to p_j and $N(\tau) \approx N(p_i)$
  $\Rightarrow \phi$ is a $(k+1)$-sliver
- $\phi \in \text{Del}(\mathcal{P})$

Inconsistency removal by weighting $\mathcal{P}$

$\omega =$ weight assignment

IF^{ω} : set of faces of the inconsistent configurations

1. For $j = 1..k$, for $i = 1..n$

   weight p_i so as to remove all slivers incident to p_i
   that are in $\mathrm{Del}^{\omega}_{T\mathbb{M}}(\mathcal{P})$ and IF^{ω}

2. for $i = 1..n$

   weight p_i so as to remove all inconsistent configurations

Basic operations

- Compute the tangent space at a point of $\mathbb{M}$
- Project a point on a k-flat
- Maintain the star of a point when varying its weight
- No d-dimensional data structure

Basic operations

- Compute the tangent space at a point of $\mathbb{M}$
- Project a point on a *k*-flat
- Maintain the star of a point when varying its weight
- No *d*-dimensional data structure

Properties of the output

If $\mathcal{P}$ is a sparse (not necessarily uniform) ε-sample, upon termination, the stars are coherent, the simplices are small and they locally approximate the tangent space of $\mathbb{M}$

- $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$ is a PL simplicial *k*-manifold
- isotopic and close to $\mathbb{M}$

Dimension of $\mathcal{S}$?

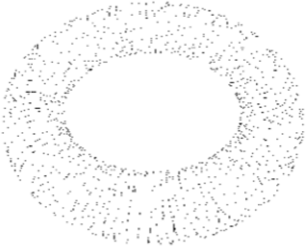

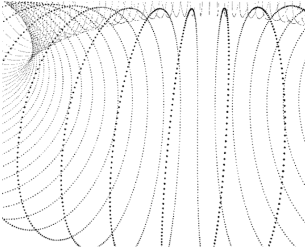

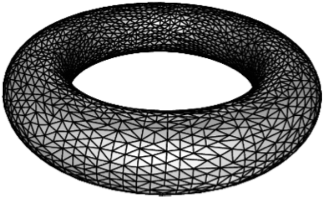

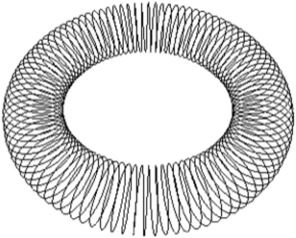

Witness complex and multiscale reconstruction

Relaxing the definition of restricted Delaunay triangulation

[Carlsson & de Silva 2004]

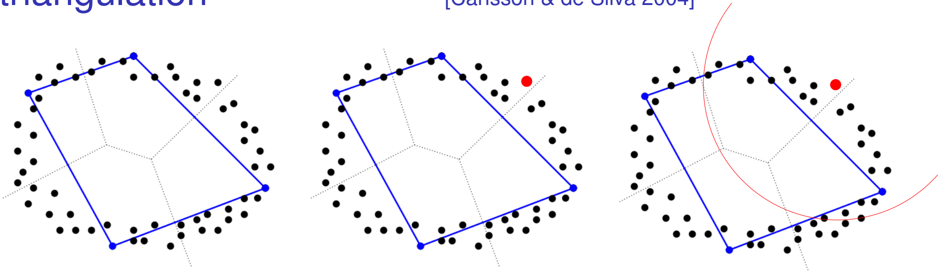

- witnesses : $W \subset \mathbb{M}$ (not necessarily finite)
- landmarks : a finite set of points $L \subset W$

- τ is a weak Delaunay simplex iff
  - $\tau \subset L$
  - $\forall \sigma \subseteq \tau$, $\exists$ $w \in W$ closer to σ than to $L \setminus \sigma$,

Weak Delaunay (witness) complex $\mathrm{Wit}(L, W)$

the collection of all weak Delaunay simplices σ,
i.e. σ and all its faces have a witness in W with respect to L

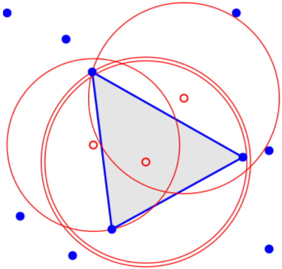

Weak Delaunay (witness) complex $\mathrm{Wit}(L, W)$

the collection of all weak Delaunay simplices σ,
i.e. σ and all its faces have a witness in W with respect to L

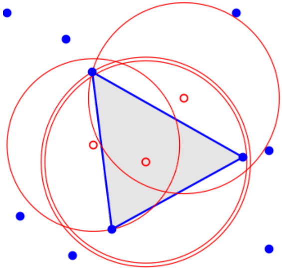

Easy to compute (only distance comparisons)

Clearly, $\mathrm{Del}(L) \subset \mathrm{Wit}(L, \mathbb{R}^d)$

Weak Delaunay (witness) complex $\mathrm{Wit}(L, W)$

the collection of all weak Delaunay simplices σ,
i.e. σ and all its faces have a witness in W with respect to L

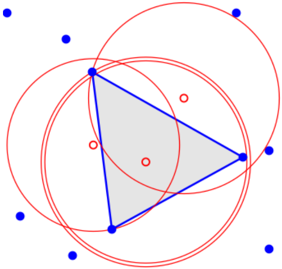

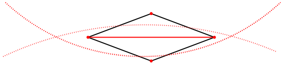

Easy to compute (only distance comparisons)

Clearly, $\mathrm{Del}(L) \subset \mathrm{Wit}(L, \mathbb{R}^d)$

Weak-witness theorem [de Silva 2008]

WW-Theorem : $\forall \Omega \subset \mathbb{R}^d, \quad \mathrm{Wit}(L, \Omega) \subset \mathrm{Wit}(L, \mathbb{R}^d) \subset \mathrm{Del}(L)$

Corollaries

- $\mathrm{Wit}(L, \mathbb{R}^d) = \mathrm{Del}(L)$
- $\mathrm{Wit}(L, \Omega)$ is embedded in $\mathbb{R}^d$

Proof of the Weak Witness theorem

$\tau = [p_0, ..., p_k]$ is a k-simplex of $\mathrm{Wit}(L)$
witnessed by a ball W_τ, i.e. $W_\tau \cap L = \tau$

We prove that $\tau \in \mathrm{Del}(L)$ by a double induction on

- k
- $|S_\tau \cap \tau| \qquad (S_\tau = \partial W_\tau)$

Clearly true for $k = 0$ and $|S_\tau \cap \tau| = k + 1$

Proof of the weak witness theorem

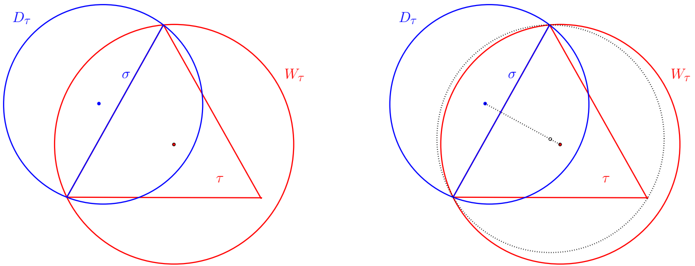

Case of sampled domains : bad news

$\mathcal{P}$ a finite set of points $\subset \Omega$, $\quad L \subset \mathcal{P}$

$\mathrm{Wit}(L, \mathcal{P}) \neq \mathrm{Del}(L, \Omega)$, even if $\mathcal{P}$ is a dense sample of Ω

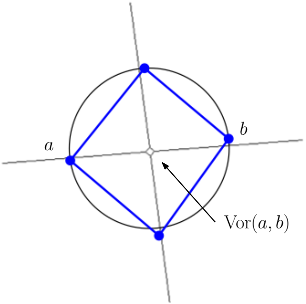

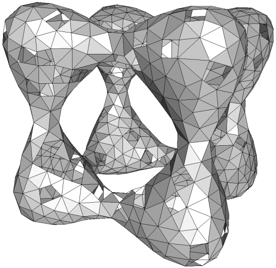

$$[ab] \in \mathrm{Wit}(L, \mathcal{P}) \quad \Leftrightarrow \quad \exists p \in \mathcal{P}, \quad \mathrm{Vor}_2(a, b) \cap \mathcal{P} \neq \emptyset$$

Case of sampled surfaces : good news

$\mathcal{P}$ a uniform ε-sample of a surface $\mathcal{S} \subset \mathbb{R}^3$

$\mathcal{P}^\xi$ the set of balls $B(p, \xi\,\mathrm{lfs}(\mathcal{S}))$, $p \in \mathcal{P}$

$L \subset \mathcal{P}$ a uniform λ-sample of $\mathcal{P}$ (landmarks)

if $\xi \geq \varepsilon$, $\mathcal{S} \subset \mathcal{P}^\xi$ and

$$\mathrm{Del}(L, \mathcal{S}) \underset{\substack{\Uparrow \\ \xi \geq \varepsilon}}{\subset} \mathrm{Wit}(L, \mathcal{P}^\xi) \underset{\substack{\Uparrow \\ \text{WW-Th}}}{\subset} \mathrm{Del}(L)$$

Although $\mathrm{Wit}(L, \mathcal{P}^\xi)$ may not be a triangulated surface, all its facets are close (both in position and orientation) to $\mathcal{S}$, which makes surface extraction easy

Case of sampled manifolds : good news

[B.,Guibas, Oudot 2008]

$\mathcal{P}$ a uniform ε-sample of a surface $\mathcal{S} \subset \mathbb{R}^3$

$\mathcal{P}^\xi$ the set of balls $B(p, \xi\, \mathrm{lfs}(\mathcal{S}))$, $p \in \mathcal{P}$

$L \subset \mathcal{P}$ a uniform λ-sample of $\mathcal{P}$ (landmarks)

$\exists\ \xi_1, \xi_2$ and a weight assignment ω s.t.

$$\begin{array}{ccccccc} \mathrm{Del}^\omega(L,\mathcal{S}) & \subset & \mathrm{Wit}^\omega(L,\mathcal{P}^{\xi_1}) & \subset & \mathrm{Cocon}^\omega(\mathcal{P}) & \subset & \mathrm{Wit}^\omega(L,\mathcal{P}^{\xi_2}) \\ & & & & \| \text{ (if no sliver)} & & \\ & & & & \mathrm{Del}^\omega(L,\mathcal{S}) & & \text{[Cheng et al. 2005]} \end{array}$$

By weighting the points of L, we can remove slivers from $\mathrm{Wit}(L, \mathcal{P}^{\xi_2})$ and obtain a triangulated manifold $\approx \mathcal{S}$

Theorem

[B., Guibas, Oudot 2007]

If $\xi_1 \geq \varepsilon$ and $30\varepsilon \leq \lambda \ll \mathrm{lfs}(\mathcal{S})$, one can compute a weight assignment ω s.t. $\mathrm{Wit}^{\omega}(L, \mathcal{P}^{\xi_1}) = \mathrm{Del}^{\omega}(L, \mathcal{S}) \approx \mathcal{S}$

Theorem

[B., Guibas, Oudot 2007]

If $\xi_1 \geq \varepsilon$ and $30\varepsilon \leq \lambda \ll \mathrm{lfs}(\mathcal{S})$, one can compute a weight assignment ω s.t. $\mathrm{Wit}^{\omega}(L, \mathcal{P}^{\xi_1}) = \mathrm{Del}^{\omega}(L, \mathcal{S}) \approx \mathcal{S}$

Multiscale reconstruction

- In practice, ε and $\mathrm{lfs}(\mathcal{S})$ are unknown
- There may exist $S_1 \ldots S_l$ s.t. W is an ε_i-sample of S_i

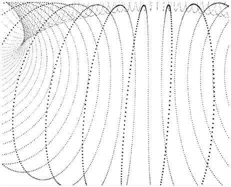

Theorem

[B., Guibas, Oudot 2007]

If $\xi_1 \geq \varepsilon$ and $30\varepsilon \leq \lambda \ll \mathrm{lfs}(\mathcal{S})$, one can compute a weight assignment ω s.t. $\mathrm{Wit}^\omega(L, \mathcal{P}^{\xi_1}) = \mathrm{Del}^\omega(L, \mathcal{S}) \approx \mathcal{S}$

Multiscale reconstruction

- In practice, ε and $\mathrm{lfs}(\mathcal{S})$ are unknown
- There may exist $S_1 ... S_l$ s.t. W is an ε_i-sample of S_i

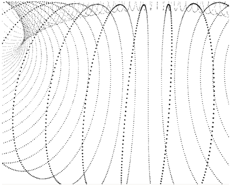

- Generate a monotonic sequence of samples $L \subseteq W$
- As long as $30\varepsilon_i \leq \lambda \leq \mathrm{lfs}(S_{i+1})$, $\mathrm{Wit}^\omega(L, \mathcal{P}^\xi) \approx \mathcal{S}_i$
- This can be detected by looking for plateaus in the diagram of Betti numbers of $\mathrm{Wit}^\omega(L, W^\xi)$

Manifold reconstruction algorithm

Greedy max min algorithm

maintain $\mathrm{Wit}^\omega(L, \mathcal{P}^{\xi_1})$ and $\mathrm{Wit}^\omega(L, \mathcal{P}^{\xi_2})$
$\xi_1 = \lambda/30, \xi_2 = 3\lambda$

INIT $L :=$ a point of $\mathcal{P}$

REPEAT $L \leftarrow p =$ the point of $\mathcal{P} \setminus L$ furthest from L
compute $\omega(p)$ so as to remove slivers in $\mathrm{Wit}^\omega(L, \mathcal{P}^{\xi_2})$
update $\mathrm{Wit}(L, \mathcal{P}^{\xi_1})$ and $\mathrm{Wit}(L, \mathcal{P}^{\xi_2})$

UNTIL $L = \mathcal{P}$

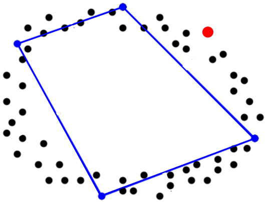

Manifold reconstruction algorithm

Greedy max min algorithm

maintain $\mathrm{Wit}^{\omega}(L, \mathcal{P}^{\xi_1})$ and $\mathrm{Wit}^{\omega}(L, \mathcal{P}^{\xi_2})$
$\xi_1 = \lambda/30, \xi_2 = 3\lambda$

INIT $L :=$ a point of $\mathcal{P}$

REPEAT $L \leftarrow p =$ the point of $\mathcal{P} \setminus L$ furthest from L
compute $\omega(p)$ so as to remove slivers in $\mathrm{Wit}^{\omega}(L, \mathcal{P}^{\xi_2})$
update $\mathrm{Wit}(L, \mathcal{P}^{\xi_1})$ and $\mathrm{Wit}(L, \mathcal{P}^{\xi_2})$

UNTIL $L = \mathcal{P}$

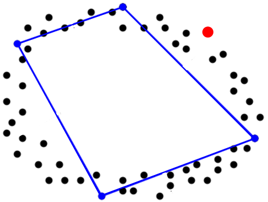

Update $\mathrm{Wit}(L, p^{\xi})$

maintain the k-order $\mathrm{Vor}(L)$, $k \leq d+1$ (The curse of dim. is back)

$$[p_1 ... p_k] \in \mathrm{Wit}(L, \mathcal{P}^{\xi}) \Leftrightarrow \exists w \in \mathcal{P}, \quad \mathrm{Vor}_k(p_1 ... p_k) \cap B(p, \xi) \neq \emptyset$$

Multi-scale reconstruction

- maintain the Betti numbers of $\hat{\mathcal{S}}$ on the fly
- detect the plateaus
- return the corresponding $\hat{\mathcal{S}}$ that are valid approximations of $\mathcal{S}$ (given the sample $\mathcal{P}$) at different scales

Multi-scale reconstruction

- maintain the Betti numbers of $\hat{\mathcal{S}}$ on the fly
- detect the plateaus
- return the corresponding $\hat{\mathcal{S}}$ that are valid approximations of $\mathcal{S}$ (given the sample $\mathcal{P}$) at different scales

Remarks

- $\mathcal{P}$ is not required to be sparse
- $\mathrm{lfs}(\mathcal{S})$ needs not to be known

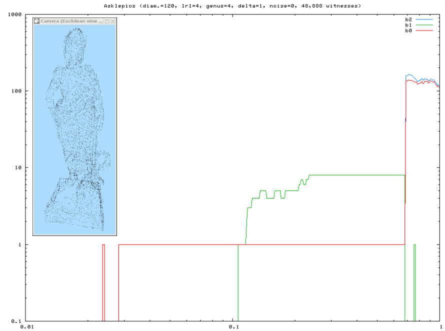

Asklepios (diam.=120, lr1=4, genus=4, delta=1, noise=0, 48,888 witnesses)
Camera (Euclidean view)
b2
b1
b0
1000
100
10
1
0.1
0.01
0.1
1

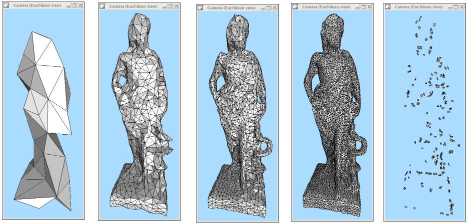

Camera (Euclidean view)
Camera (Euclidean view)
Camera (Euclidean view)
Camera (Euclidean view)
Camera (Euclidean view)

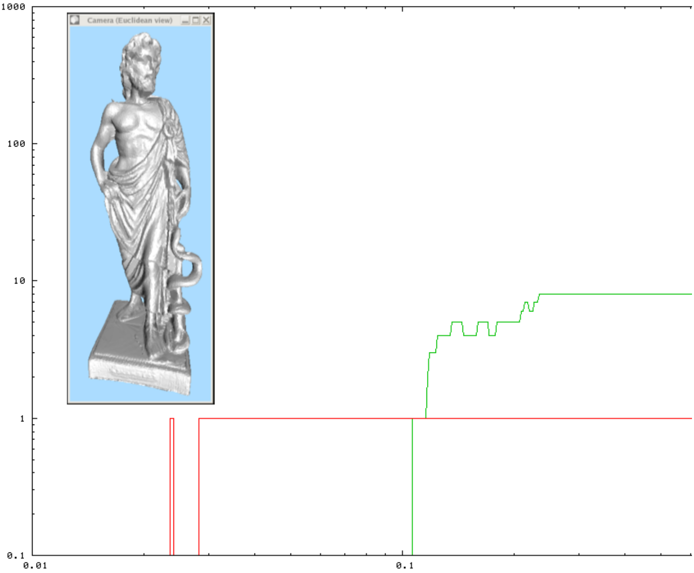
Camera (Euclidean view)
1000
100
10
1
0.1
0.01
0.1

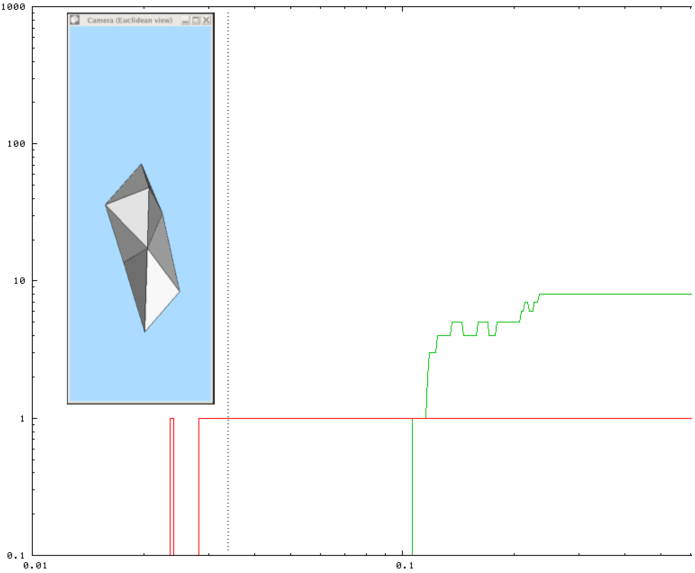
Camera (Euclidean view)
1000
100
10
1
0.1
0.01
0.1

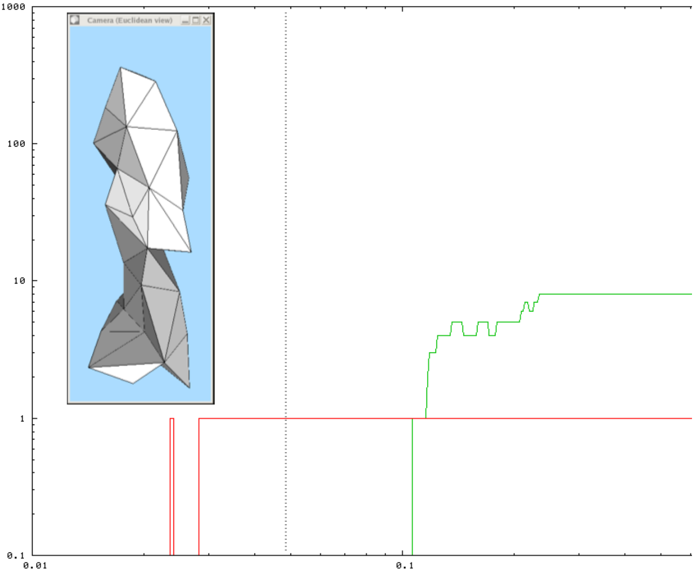
1000
100
10
1
0.1
0.01
0.1
Camera (Euclidean view)

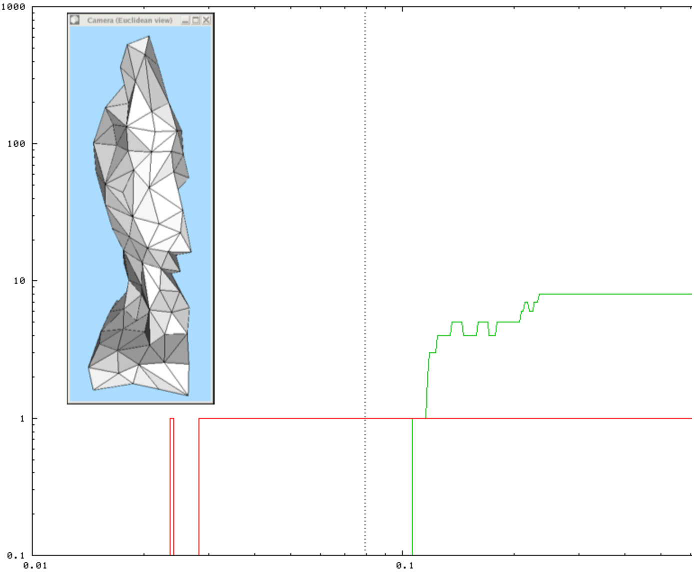

1000
100
10
1
0.1
0.01
0.1
Camera (Euclidean view)

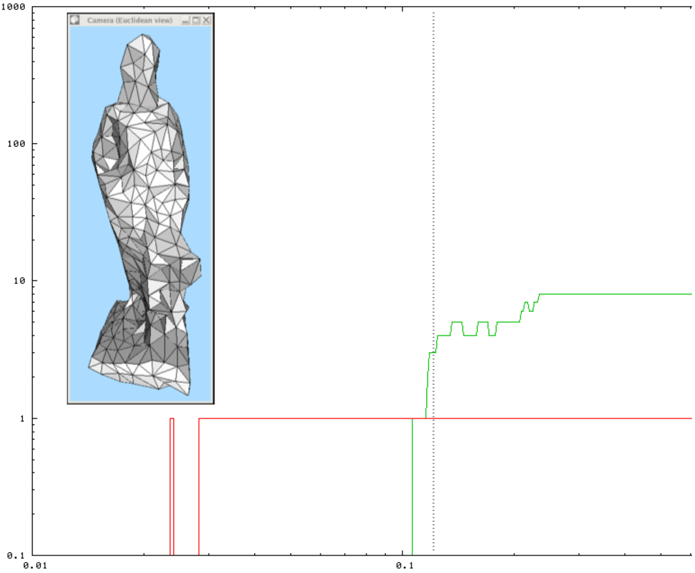
Camera (Euclidean view)
1000
100
10
1
0.1
0.01
0.1

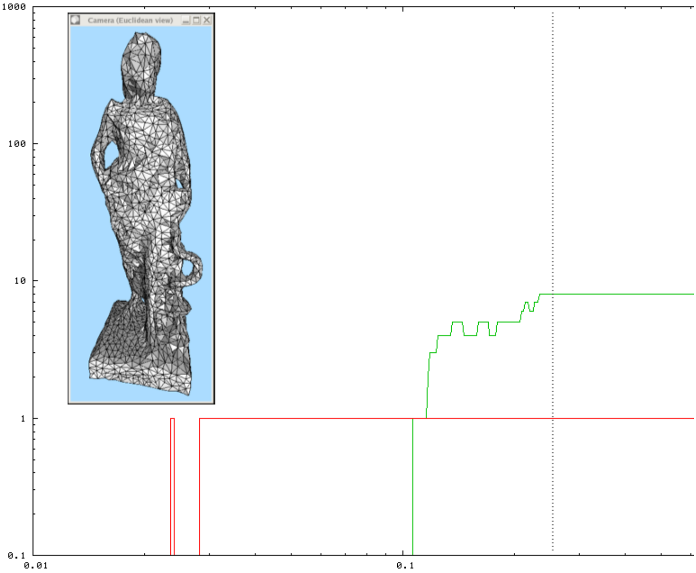
1000
100
10
1
0.1
0.01
0.1
Camera (Euclidean view)

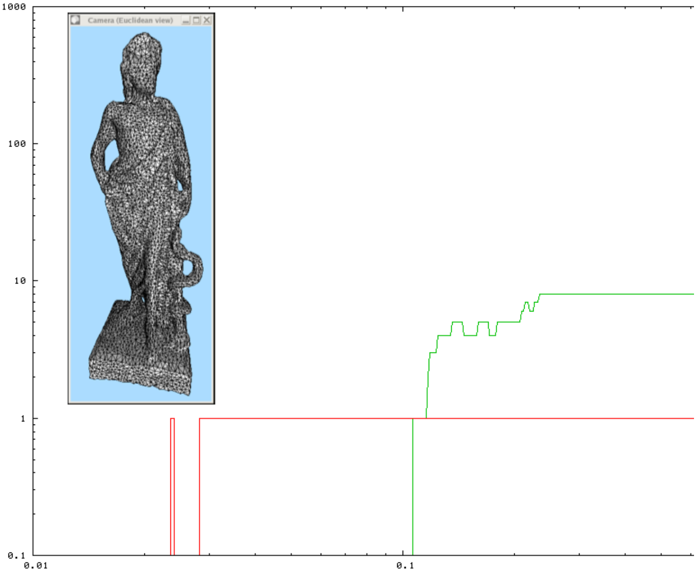
1000
100
10
1
0.1
0.01
0.1
Camera (Euclidean view)