Estimating Geometry and Topology from Voronoi Diagrams

Tamal K. Dey

The Ohio State University

Voronoi diagrams

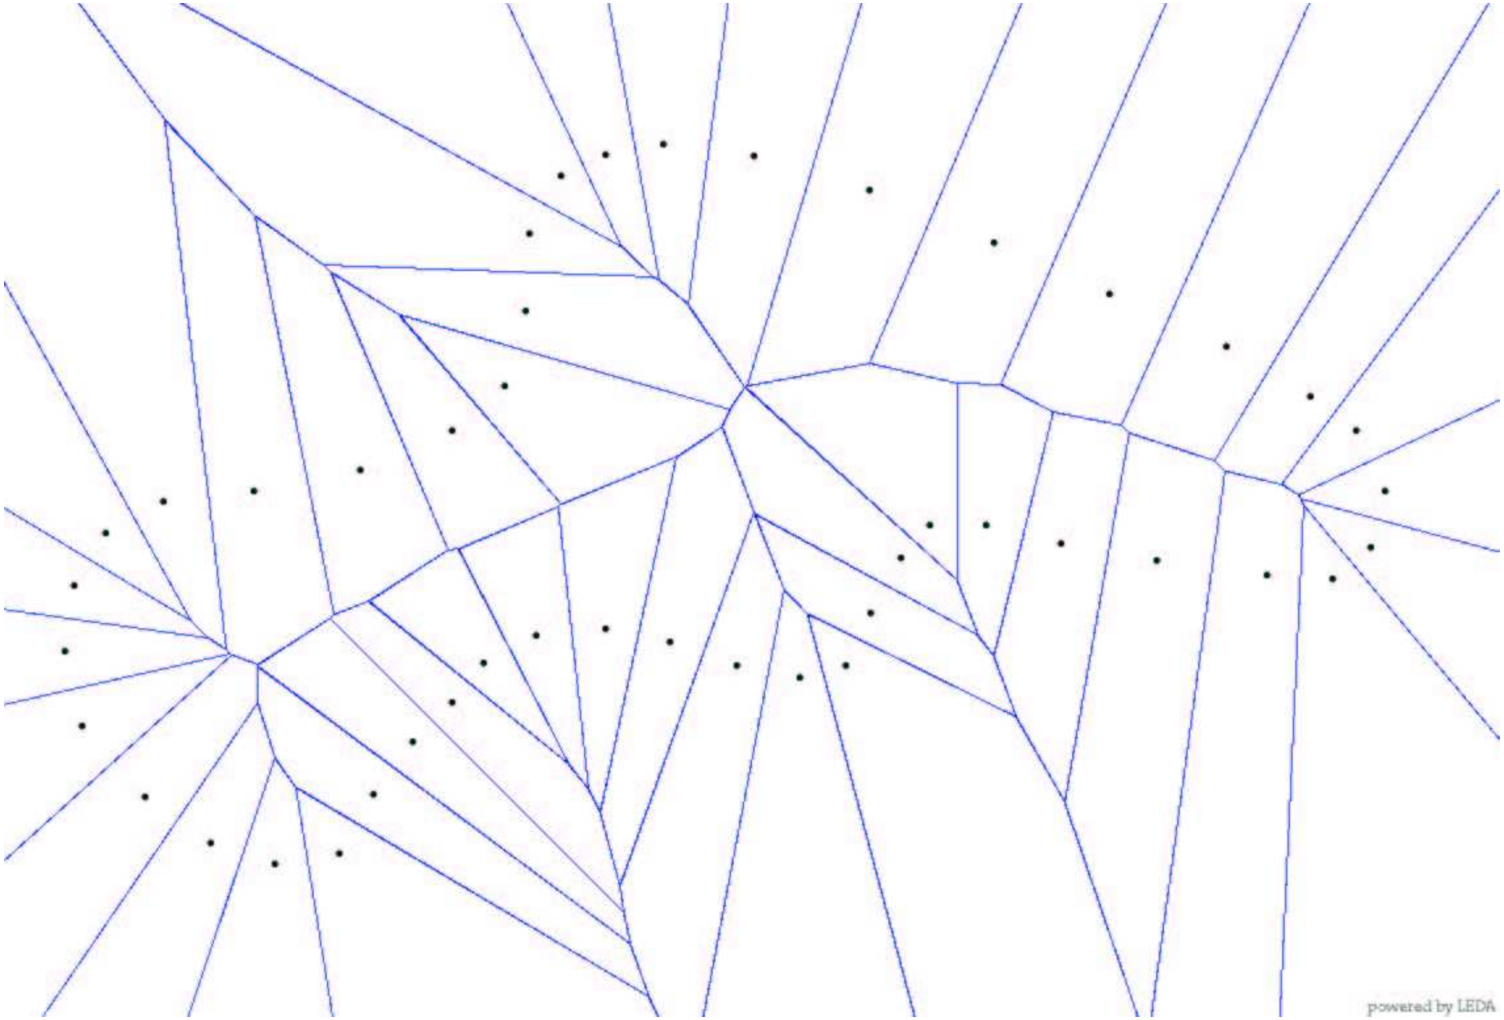

Voronoi diagrams

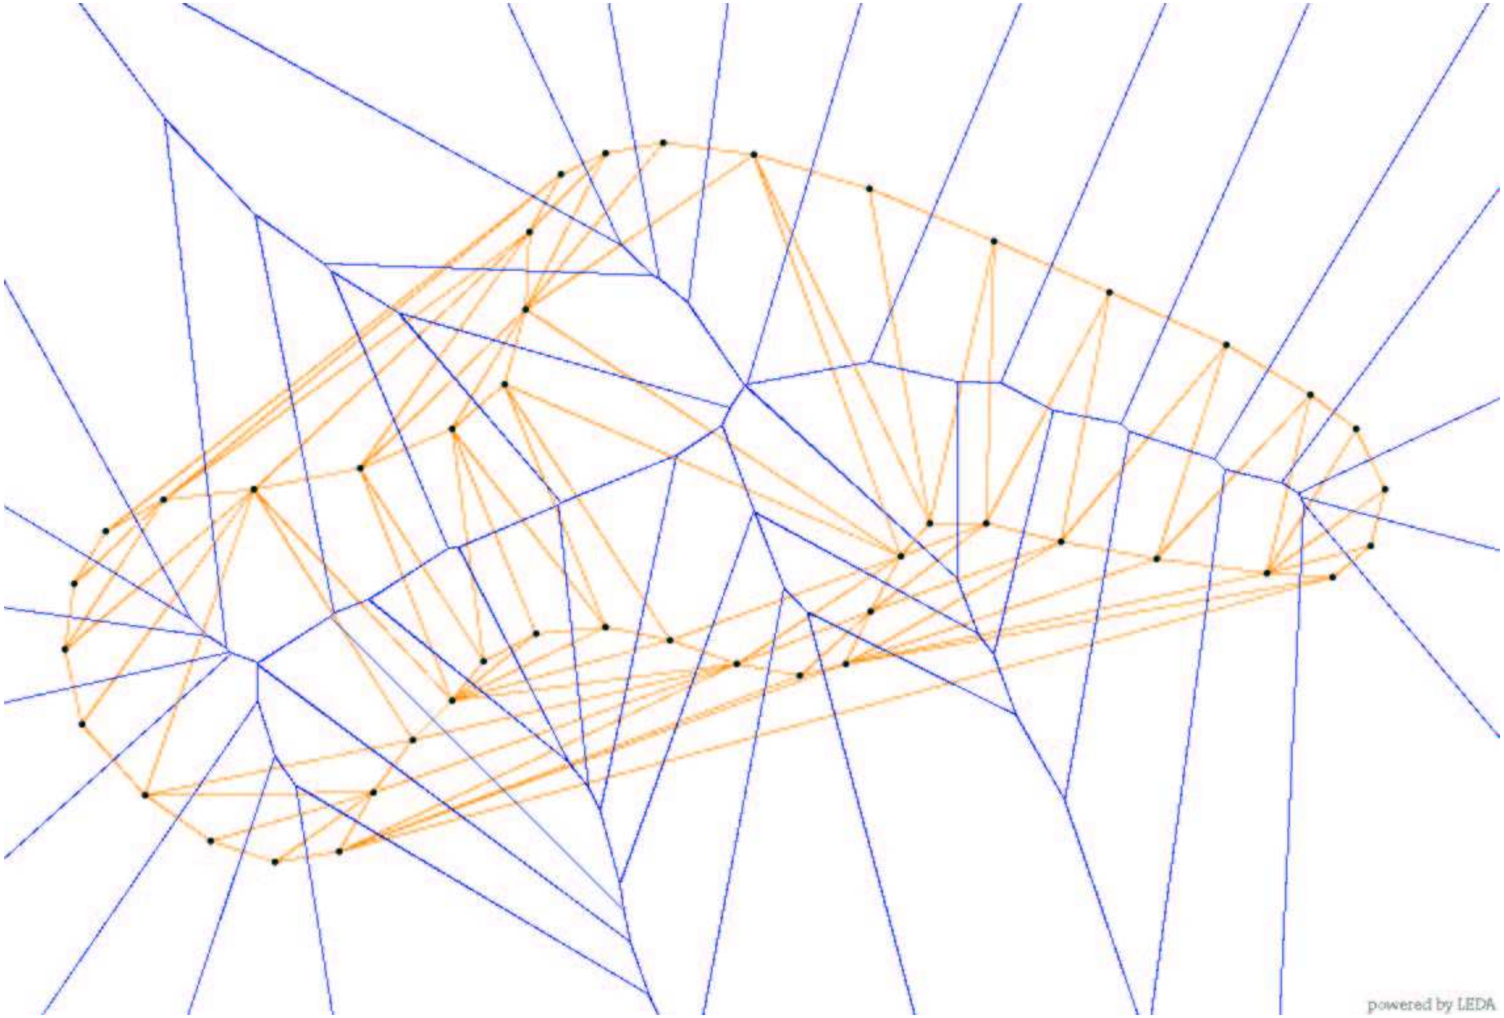

Voronoi diagrams

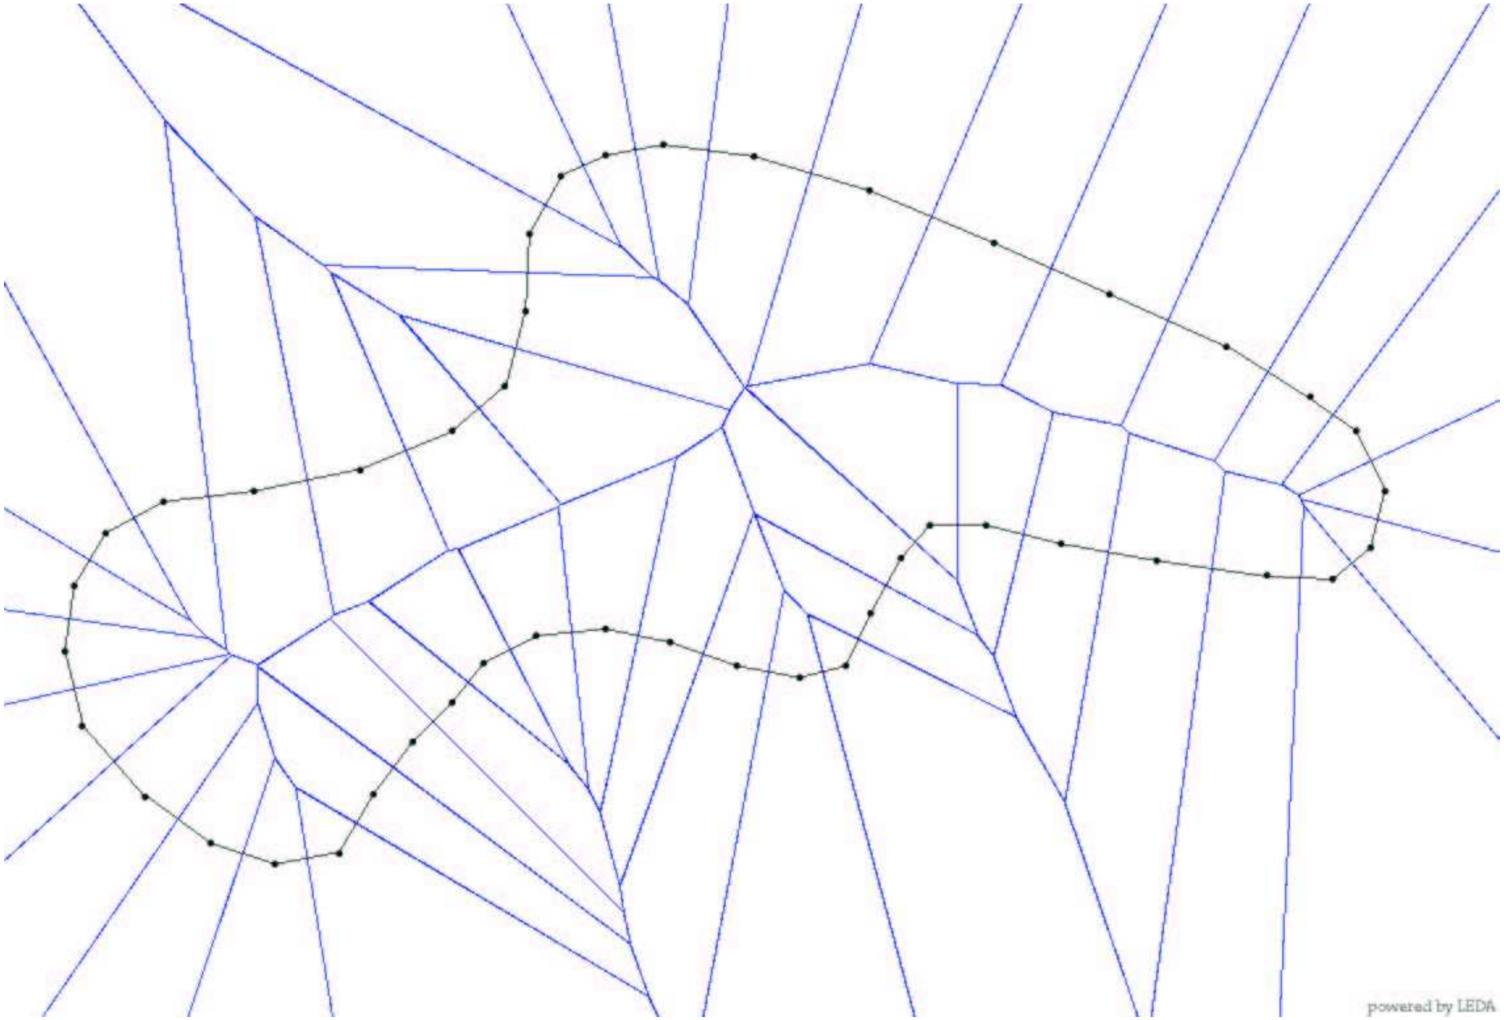

Problem

Estimating geometry: Given P presumably sampled from a k-dimensional manifold $\mathcal{M} \subset \mathbb{R}^d$ estimate geometric attributes such as normals, curvatures of $\mathcal{M}$ from $\mathrm{Vor}\, P$.

Estimating topology: (i) capture the major topological features (persistent topology) of $\mathcal{M}$ from $\mathrm{Vor}\, P$ (ii) capture the exact topology of $\mathcal{M}$ from $\mathrm{Vor}\, P$.

Three dimensions

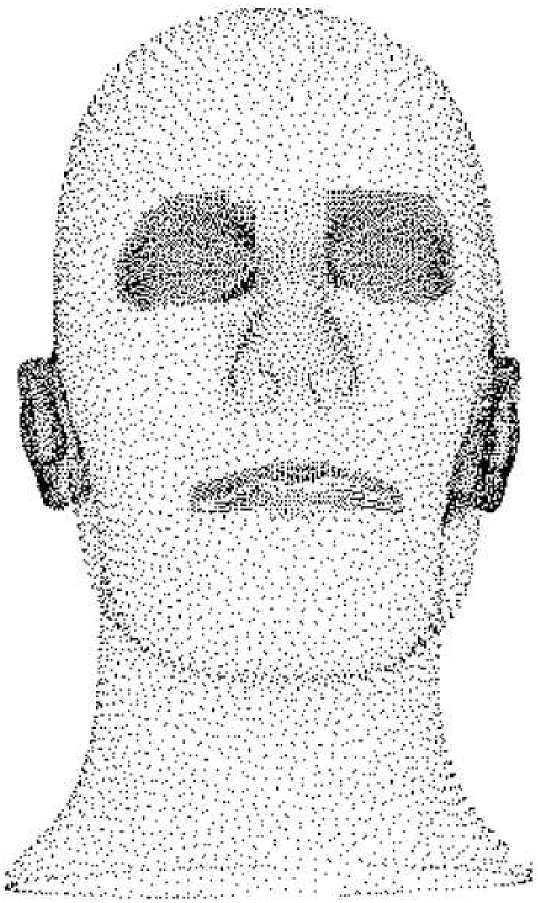

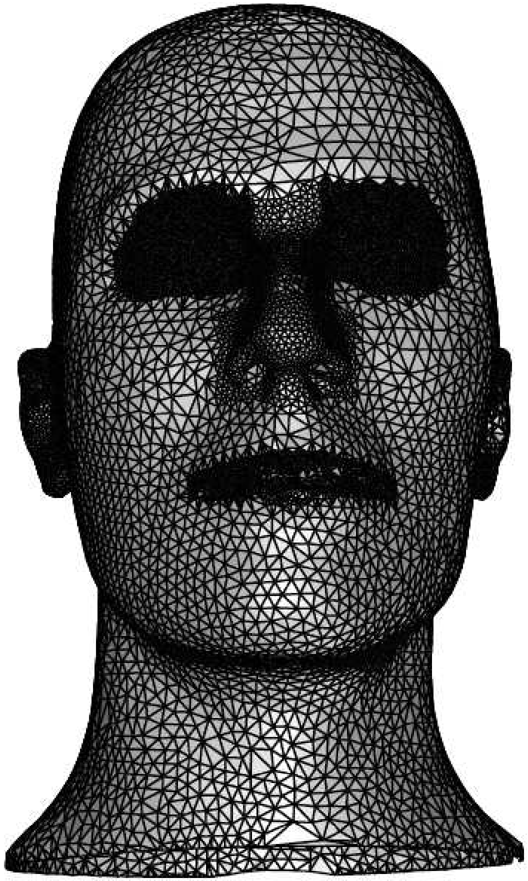

Medial Axis and Local Feature Size

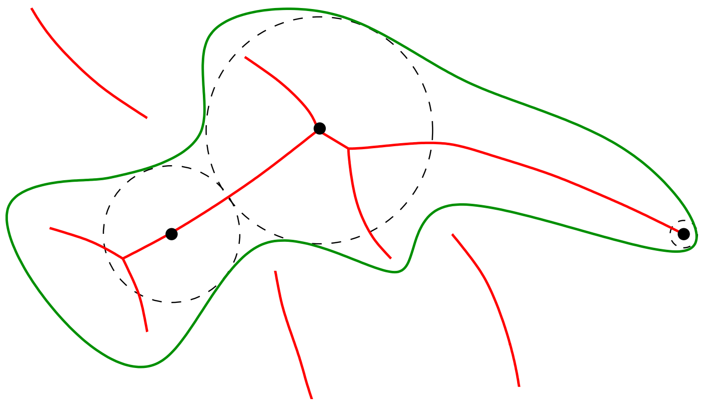

MEDIAL AXIS:

Set of centers of maximal empty balls.

LOCAL FEATURE SIZE:

For $x \in \mathcal{M}$, $f(x)$ is the distance to the medial axis.

$f(x) \leq f(y) + \|xy\|$ 1-Lipschitz

Good Sampling

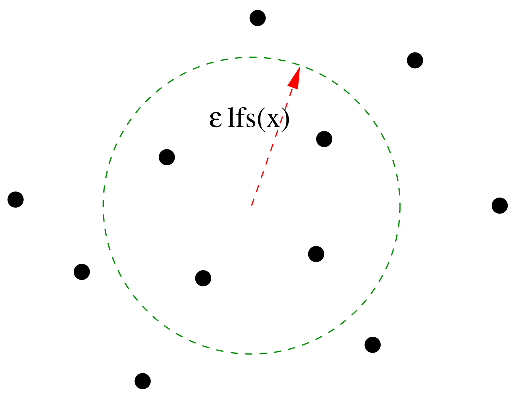

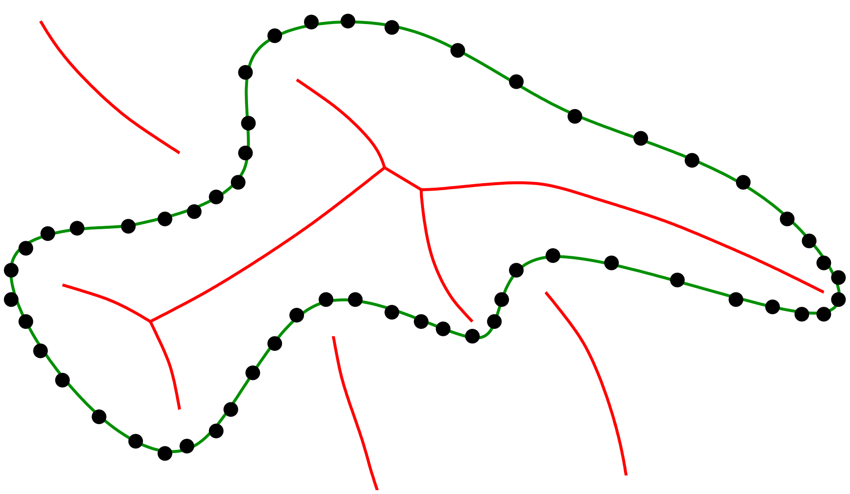

ε-SAMPLING[AMENTA-BERN-EPPSTEIN 97]: $P \subset \mathcal{M}$ such that

$$\forall x \in \mathcal{M}, \quad B(x, \varepsilon \cdot f(x)) \cap P \neq \emptyset.$$

Normal estimation

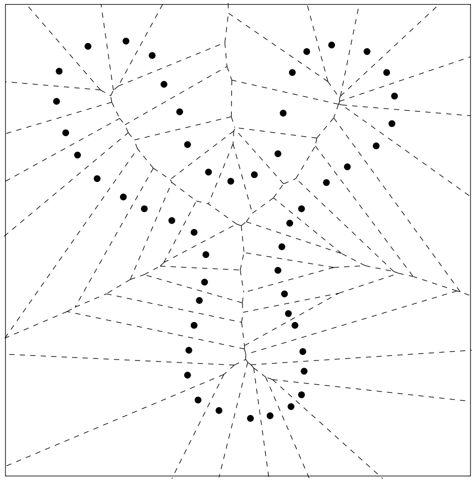

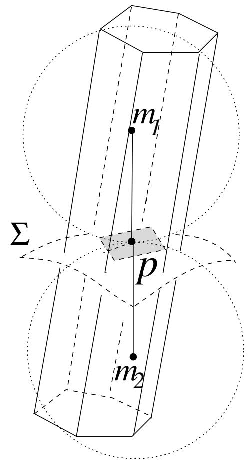

Normal estimation

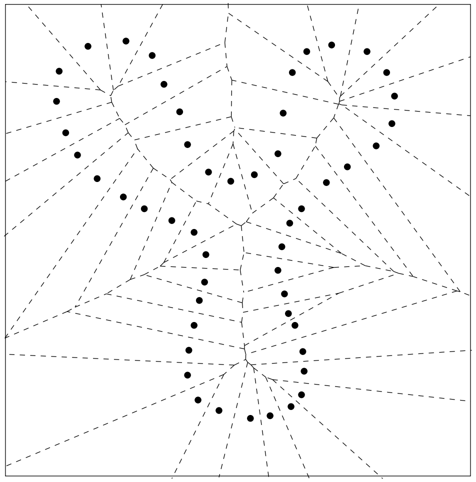

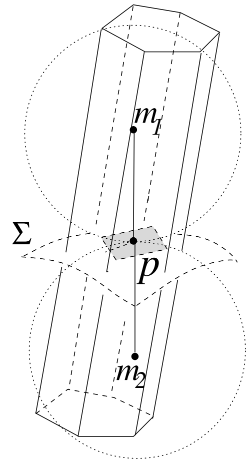

Normal Lemma [Amenta-Bern 98] : For $\varepsilon < 1$, the angle (acute) between the normal $\mathbf{n}_p$ at p and the pole vector $\mathbf{v}_p$ is at most

$$2 \arcsin \frac{\varepsilon}{1-\varepsilon}.$$

Noise

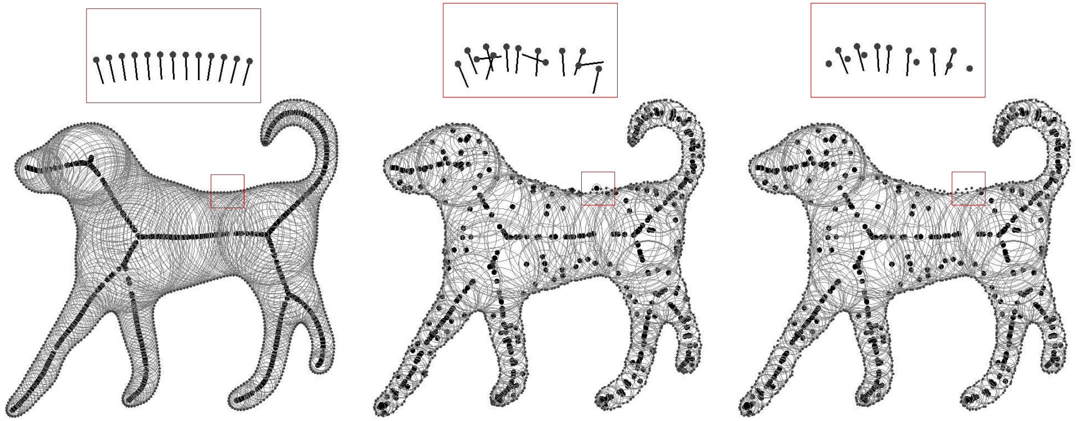

Noise

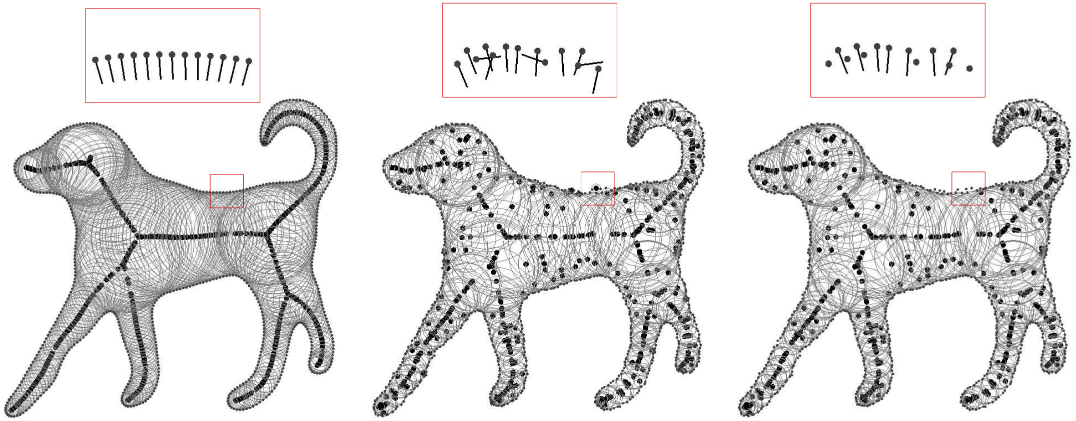

- $\tilde{p}$, $\tilde{P}$ are orthogonal projections of p and P on $\mathcal{M}$.

Noise

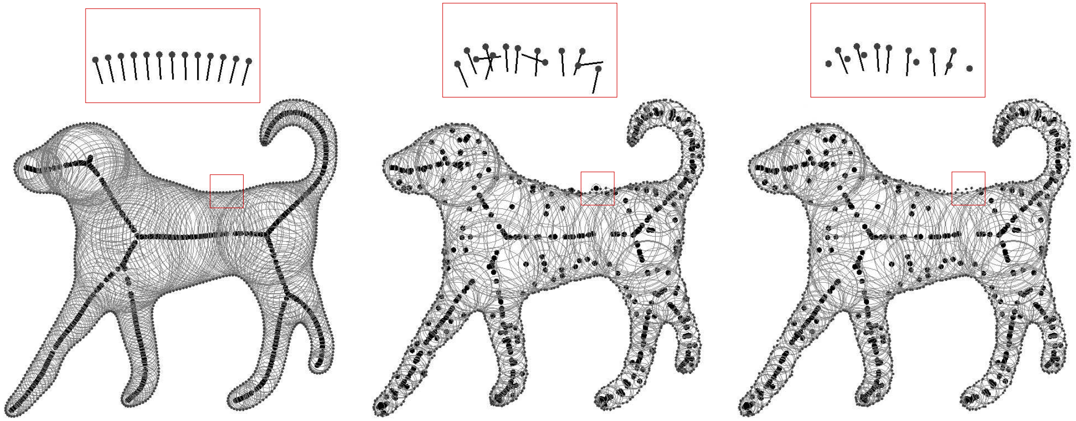

- $\tilde{p}$, $\tilde{P}$ are orthogonal projections of p and P on $\mathcal{M}$.

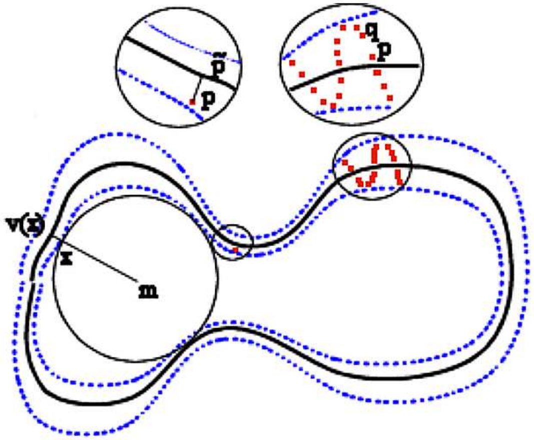

Noise

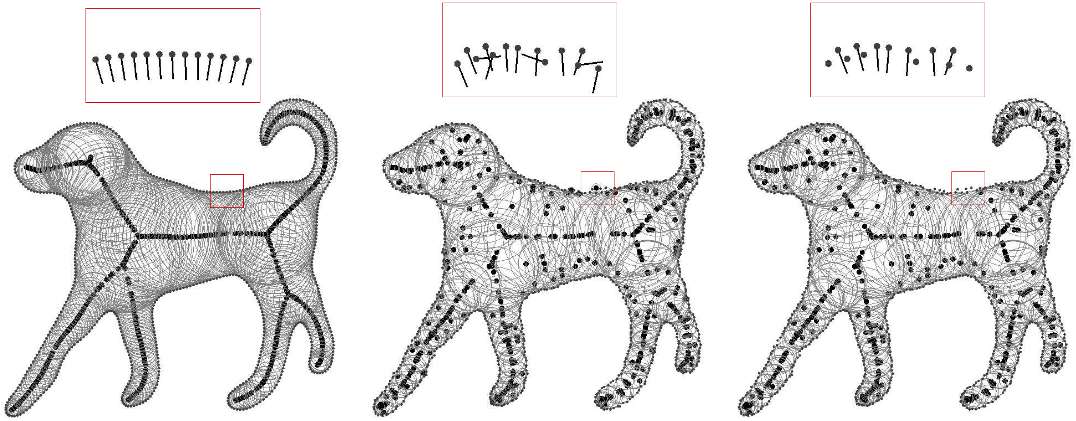

- $\tilde{p}$, $\tilde{P}$ are orthogonal projections of p and P on $\mathcal{M}$.
- P is (ε, η)-sample of $\mathcal{M}$ if
  - $\tilde{P}$ is a ε-sample of $\mathcal{M}$,
  - $d(p, \tilde{p}) \leq \eta f(\tilde{p})$.

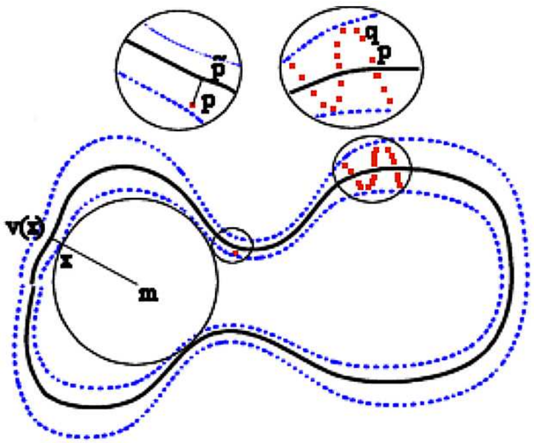

Noise and normals

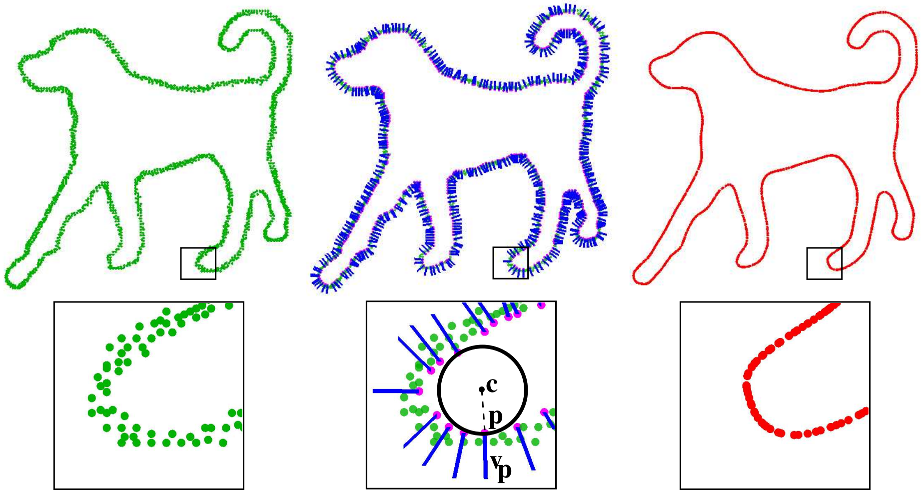

Noise and normals

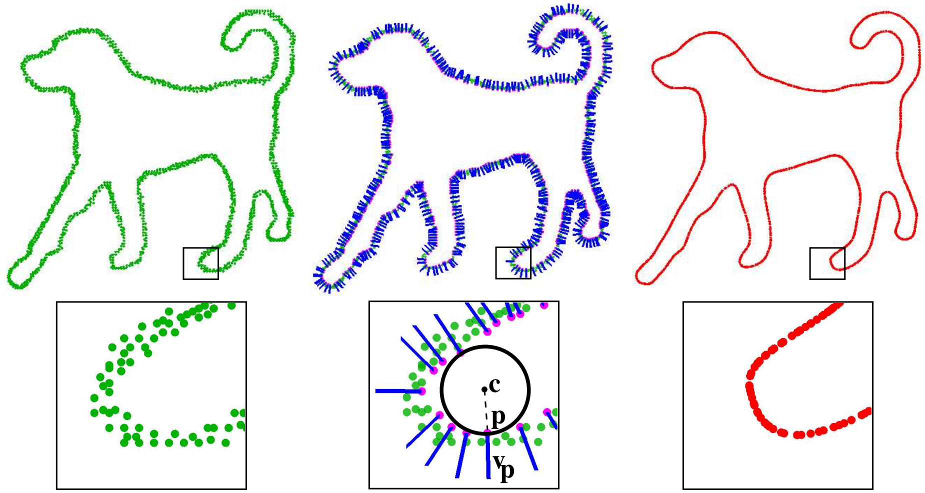

Normal Lemma [Dey-Sun 06] : Let $p \in P$ with $d(p, \tilde{p}) \leq \eta f(\tilde{x})$ and $B_{c,r}$ be any Delaunay/Voronoi ball incident to p so that $r = \lambda f(\tilde{p})$). Then,

$$\angle cx, \mathbf{n}_{\tilde{x}} = O(\frac{\varepsilon}{\lambda} + \sqrt{\frac{\eta}{\lambda}}).$$

Noise and normals

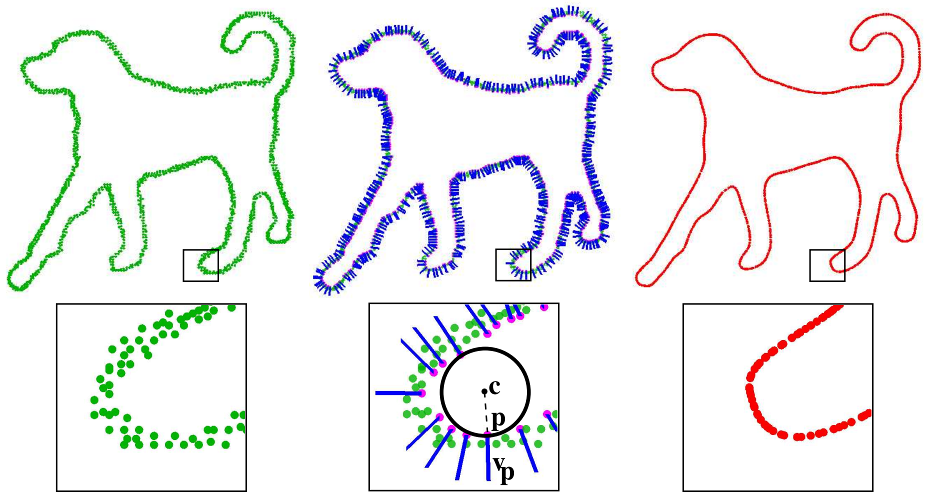

Normal Lemma [Dey-Sun 06] : Let $p \in P$ with $d(p, \tilde{p}) \leq \eta f(\tilde{x})$ and $B_{c,r}$ be any Delaunay/Voronoi ball incident to p so that $r = \lambda f(\tilde{p})$). Then,

$$\angle cx, \mathbf{n}_{\tilde{x}} = O(\frac{\varepsilon}{\lambda} + \sqrt{\frac{\eta}{\lambda}}).$$

- Gives an algorithm to estimate normals.

Noise and features

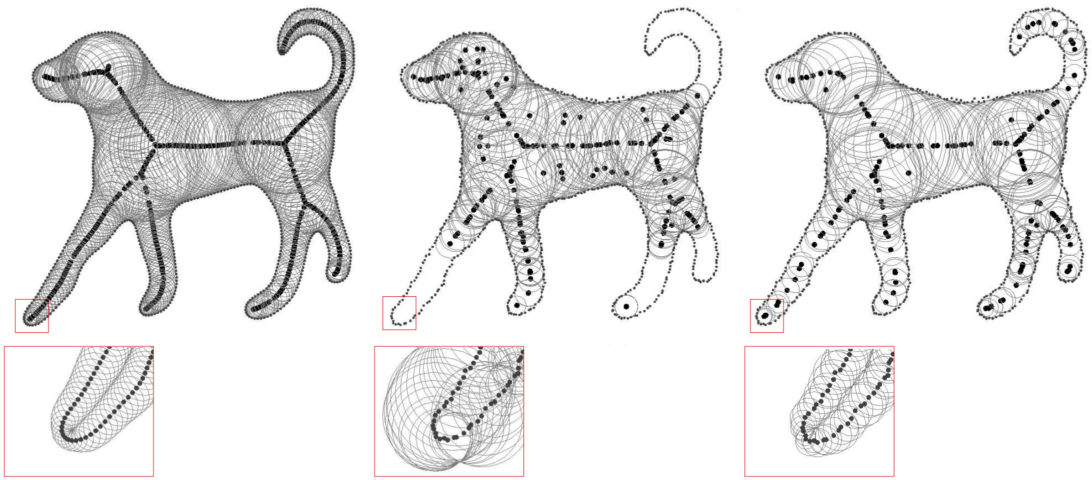

Noise and features

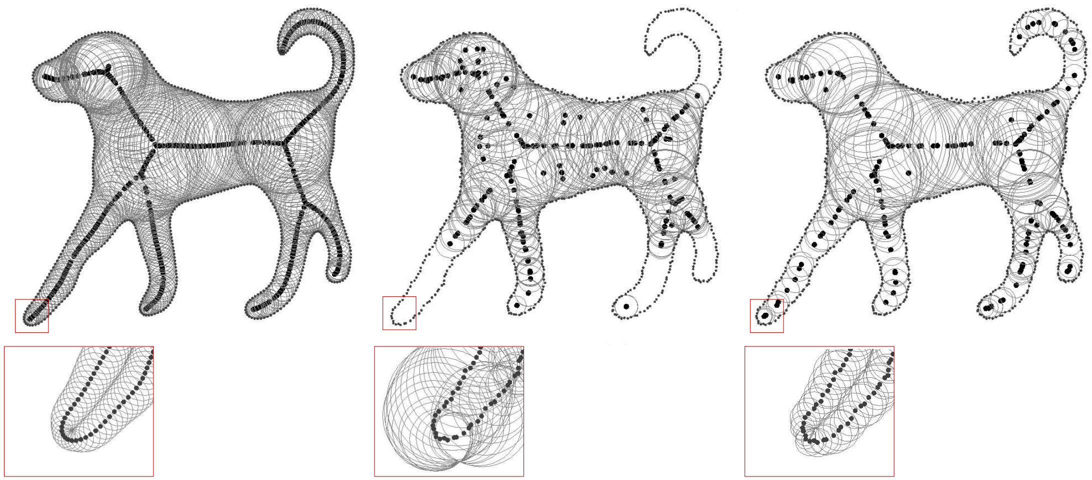

Medial Axis Approximation Lemma [Dey-Sun 06] : If P is a (ε, η)-sample of $\mathcal{M}$, then the medial axis of $\mathcal{M}$ can be approximated with Hausdorff distance of $O(\varepsilon^{\frac{1}{4}} + \eta^{\frac{1}{4}})$ times the respective medial ball radii.

Noise and features

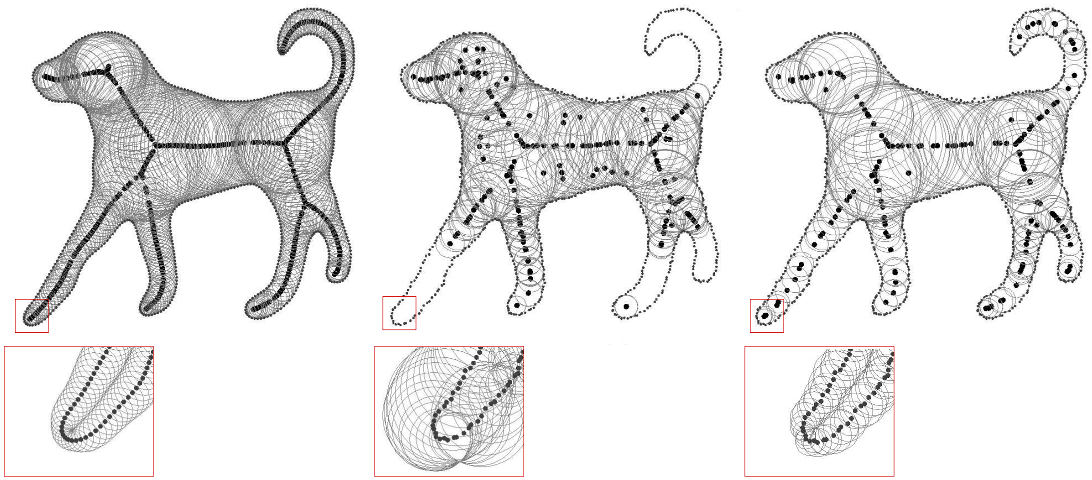

Medial Axis Approximation Lemma [Dey-Sun 06] : If P is a (ε, η)-sample of $\mathcal{M}$, then the medial axis of $\mathcal{M}$ can be approximated with Hausdorff distance of $O(\varepsilon^{\frac{1}{4}} + \eta^{\frac{1}{4}})$ times the respective medial ball radii.

- Gives an algorithm to estimate local feature size.

Topology

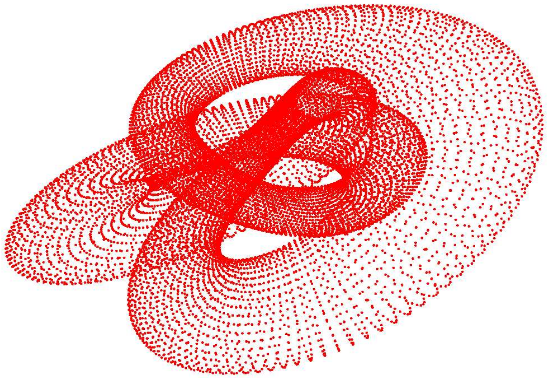

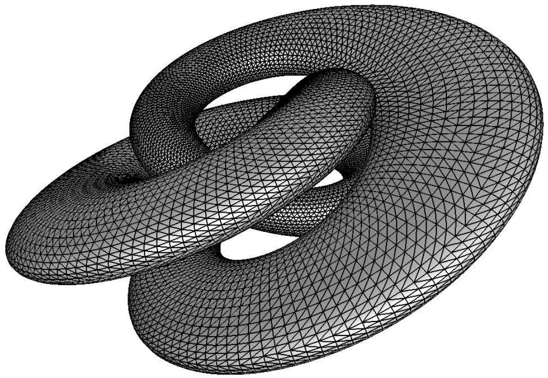

Topology

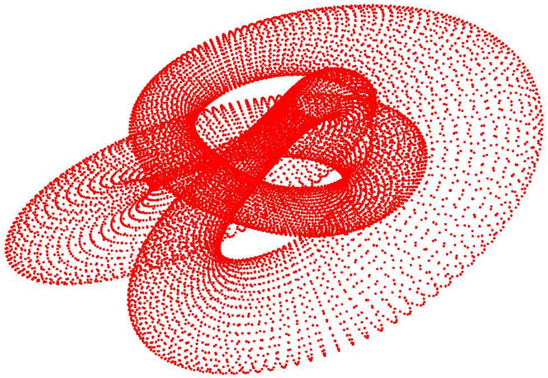 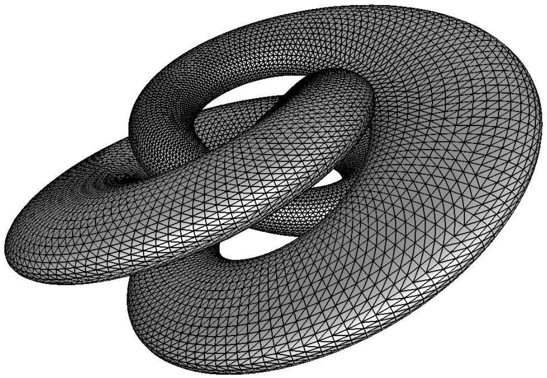

Homeomorphic/Isotopic reconstruction [ACDL00]: Let $P \subset \mathcal{M}$ be ε-sample. A Delaunay mesh $T \subset \mathrm{Del}\, P$ can be computed so that

Topology

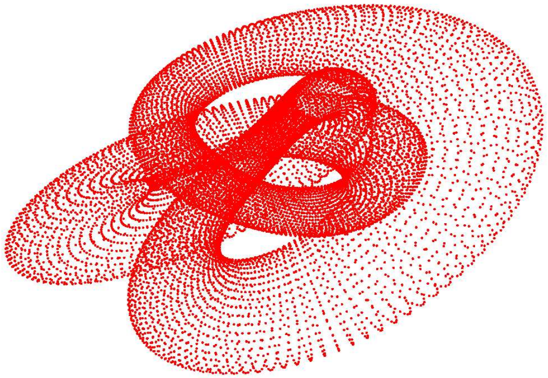

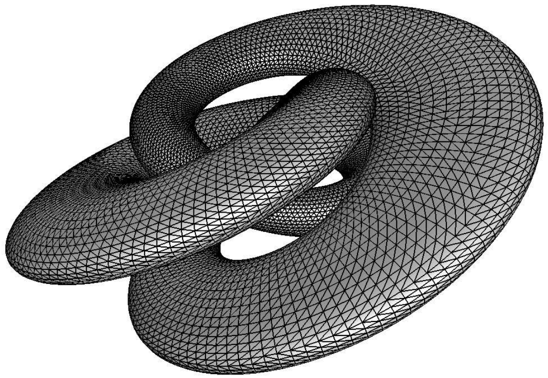

Homeomorphic/Isotopic reconstruction [ACDL00]: Let $P \subset \mathcal{M}$ be ε-sample. A Delaunay mesh $T \subset \mathrm{Del}\, P$ can be computed so that

- there is an isotopy $h : |T| \times [0,1] \to \mathbb{R}^3$ between $|T|$ and $\mathcal{M}$. Moreover, $h(|T|, 1)$ is the orthogonal projection map.
- the isotopy moves any point $x \in |T|$ only by $O(\varepsilon) f(\tilde{x})$ distance.
- triangles in T have normals within $O(\varepsilon)$ angle of the respective normals at the vertices.

Original Crust algorithm [AB98], Cocone algorithm [ACDL00], Natural neighbor algorithm [BC00] enjoy these properties.

Cocone Algorithm

Amenta, Choi, Dey and Leekha

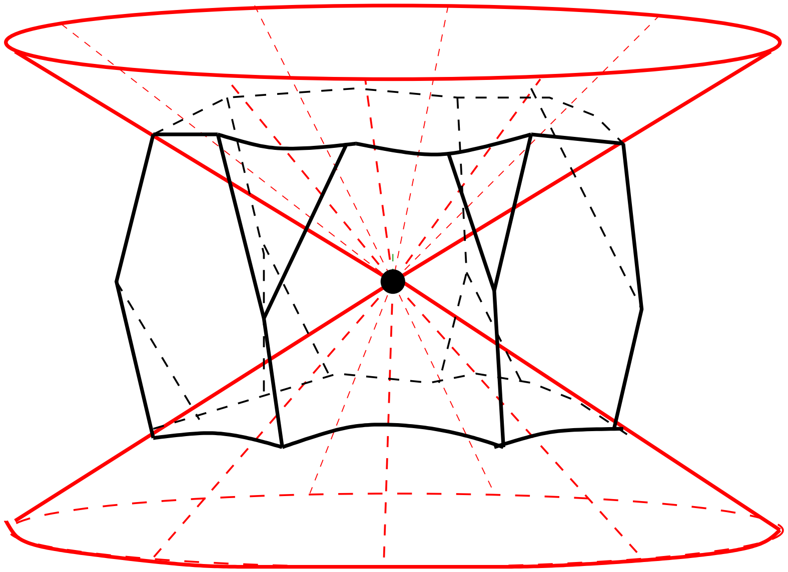

Cocone Algorithm

Amenta, Choi, Dey and Leekha

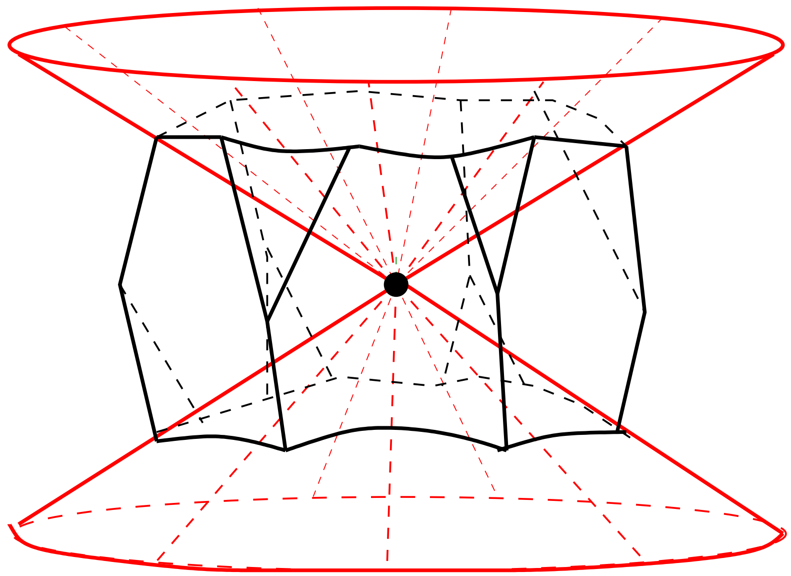

Cocone triangles for p: Delaunay triangles incident to p that are dual to Voronoi edges inside the cocone region.

Cocone Algorithm

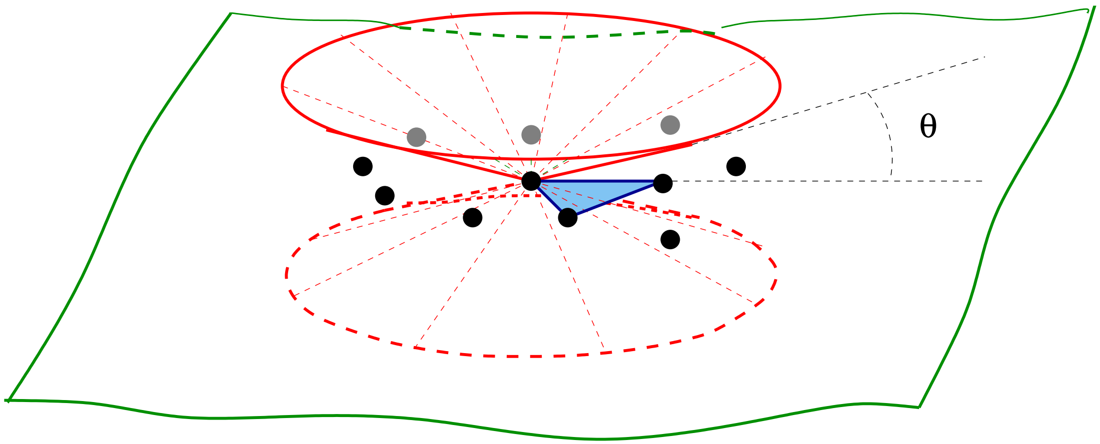

Cocone Algorithm

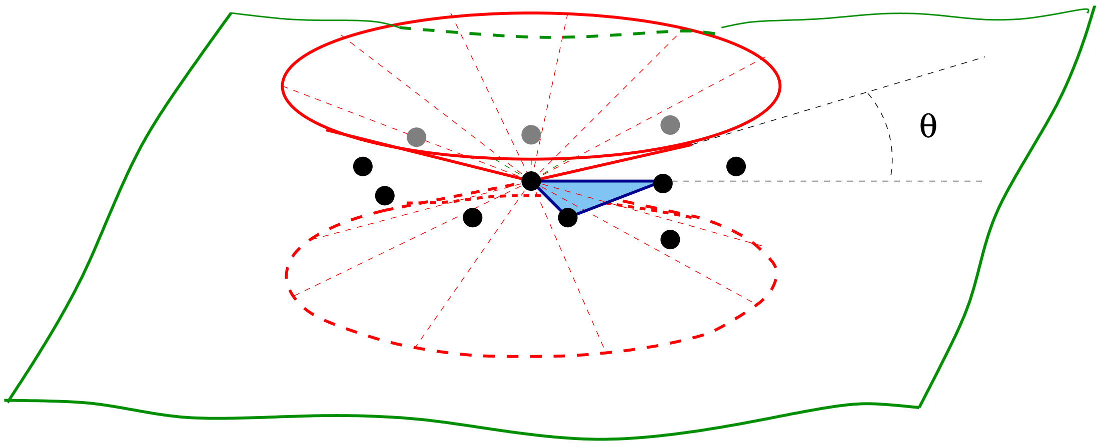

Cocone triangles :

- are nearly orthogonal to the estimated normal at p
- have empty spheres that are near equatorial.

Noisy sample : Topology

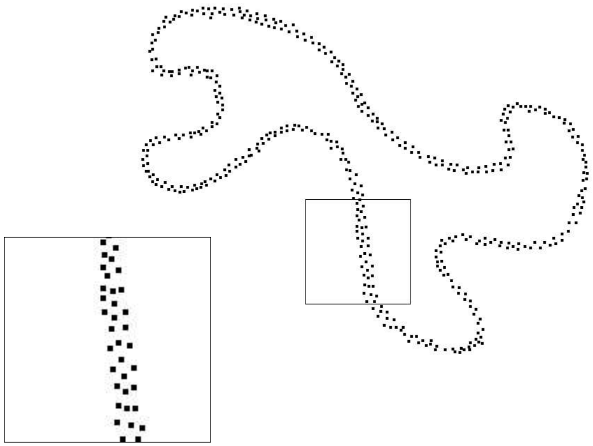

Input noisy sample

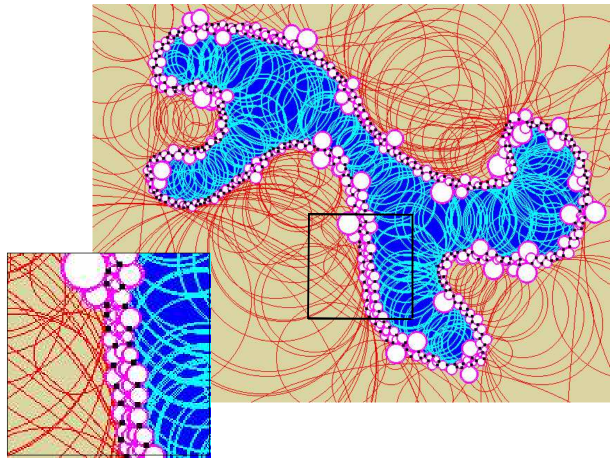

Step 1

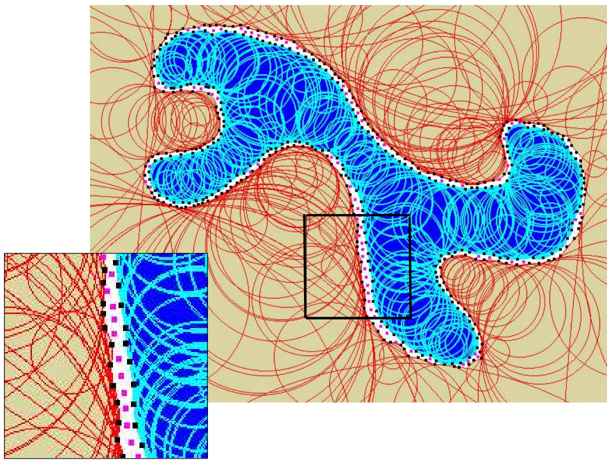

Step 2

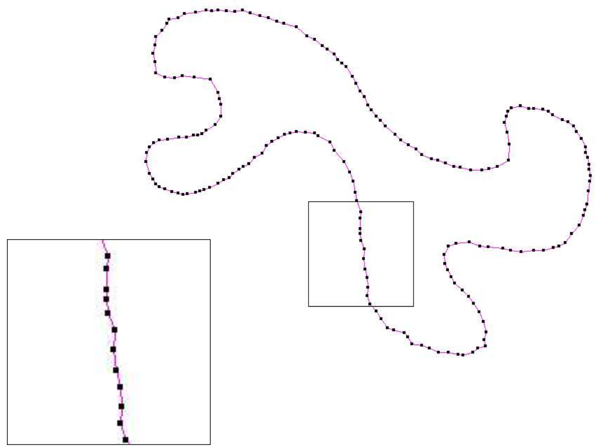

Step 3

Noisy sample : topology

Homeomorphic reconstruction [Dey-Goswami 04] : Let $P \subset \mathcal{M}$ be $(\varepsilon, \varepsilon^2)$-sample. For $\lambda > 0$, Let $B_\lambda = \{B_{c,r}\}$ be the set of (inner) Delaunay balls where $r > \lambda f(\tilde{c})$. There exists a $\lambda > 0$ so that $\mathrm{bd} \bigcup B_\lambda \approx \mathcal{M}$.

Noisy sample : topology

Homeomorphic reconstruction [Dey-Goswami 04] : Let $P \subset \mathcal{M}$ be $(\varepsilon, \varepsilon^2)$-sample. For $\lambda > 0$, Let $B_\lambda = \{B_{c,r}\}$ be the set of (inner) Delaunay balls where $r > \lambda f(\tilde{c})$. There exists a $\lambda > 0$ so that $\mathrm{bd} \bigcup B_\lambda \approx \mathcal{M}$.

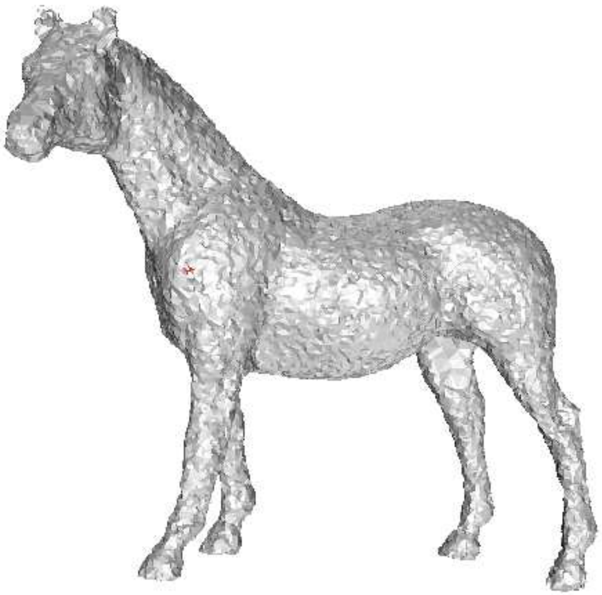

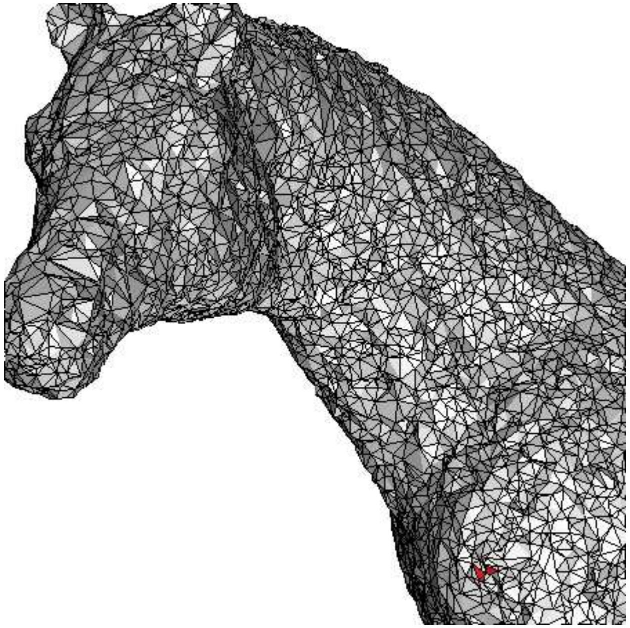

Higher dimensions

Assumptions: Sample P from a smooth, compact k-manifold $\mathcal{M} \subset \mathbb{R}^d$ without boundary. P is "sufficiently dense and uniform".

Higher dimensions

Assumptions: Sample P from a smooth, compact k-manifold $\mathcal{M} \subset \mathbb{R}^d$ without boundary. P is "sufficiently dense and uniform".

- Normal space estimation, dimenison detection;
- Homeomorphic reconstruction

Good Sampling

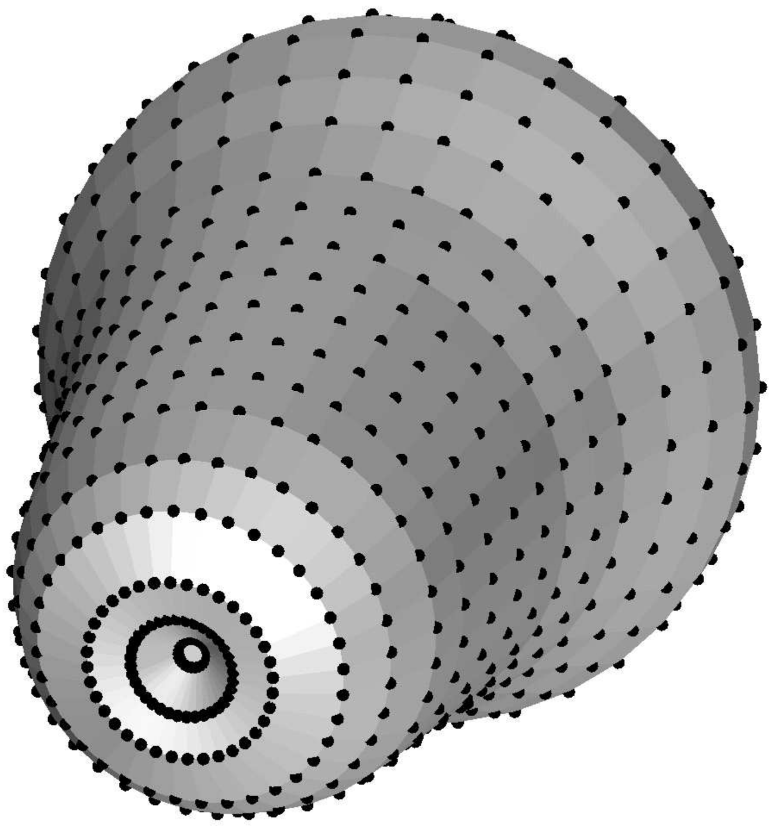

Good Sampling

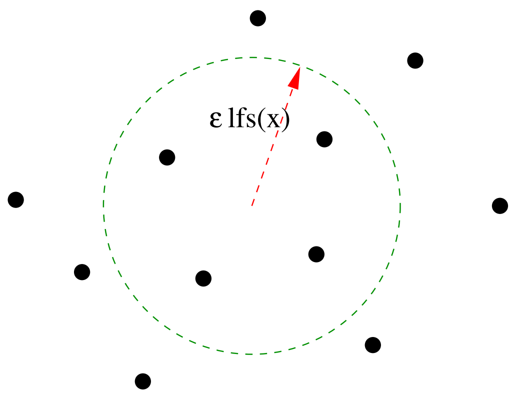

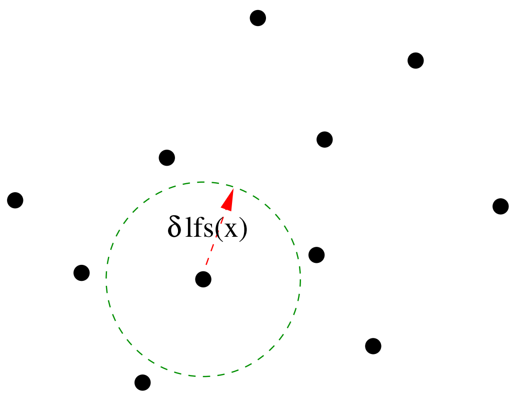

$0 < \delta < \varepsilon$

(ε, δ)-SAMPLING: $P \subset \mathcal{M}$ such that

$\forall x \in \mathcal{M}, \quad B(x, \varepsilon \cdot f(x)) \cap P \neq \emptyset.$

$\forall p \in P, \quad B(p, \delta \cdot f(p)) \cap P = \{p\}.$

Dimension detection

Dey-Giesen-Goswami-Zhao 2002

Define Voronoi subset V_p^i recursively for a point $p \in P$.

- $\text{aff}\, V_p^k$ approximates T_p.
- k can be determined if P is (ε, δ)-sample for appropriate ε and δ.

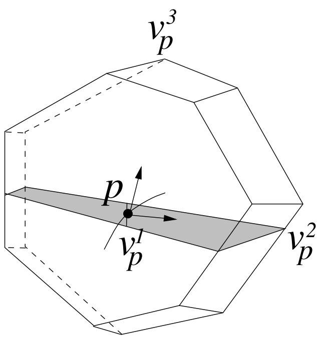

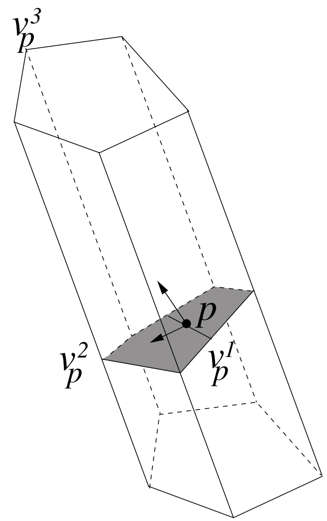

Restricted Voronoi Diagram

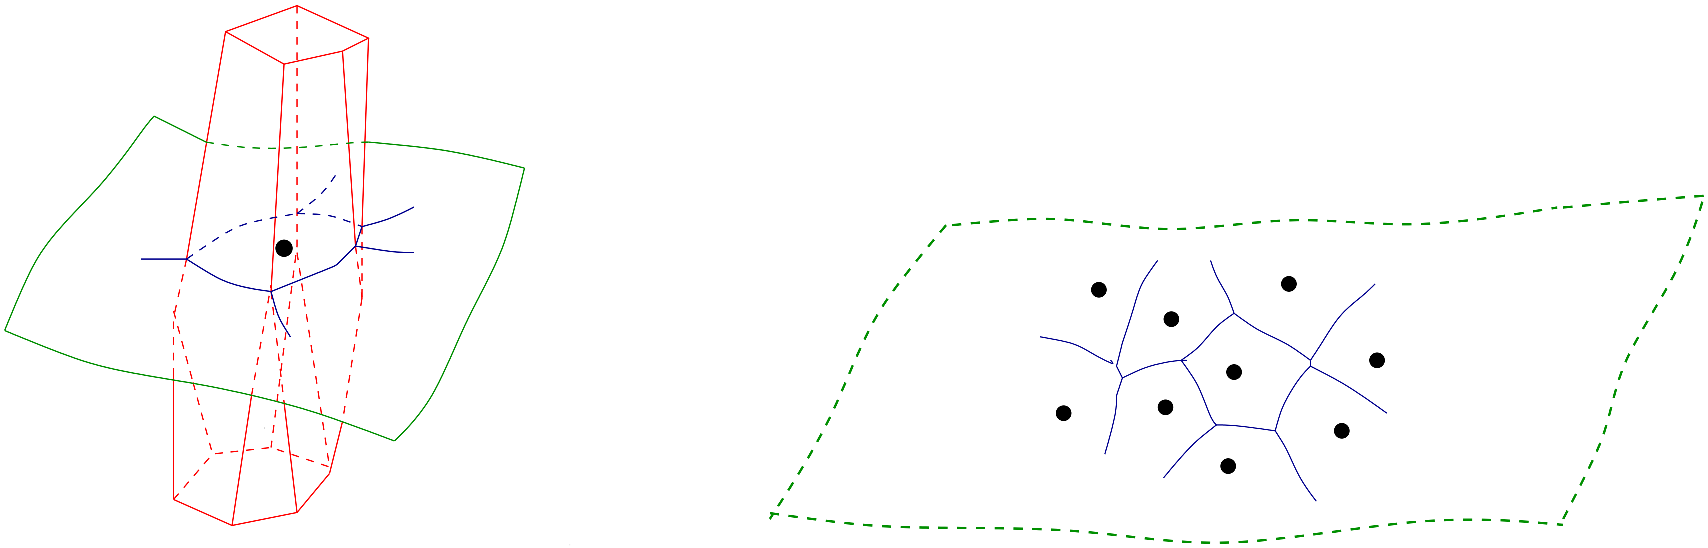

RESTRICTED VORONOI FACE: Intersection of Voronoi face with manifold

BALL PROPERTY: Each Voronoi face is topologically a ball.

Restricted Delaunay Triangulation

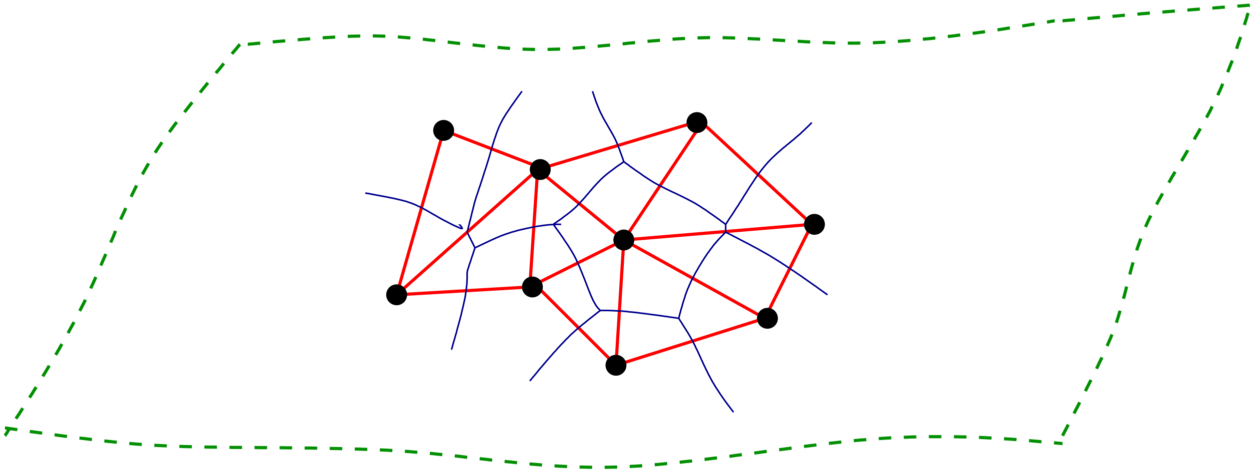

$\mathrm{Del}_{\mathcal{M}}\, P$

This is a good candidate to be a “correct reconstruction”.

Topology of RDT

TBP Theorem [Edelsbrunner-Shah 94] : If $\mathrm{Vor}\, P$ has the topological ball property w.r.t. $\mathcal{M}$, then $\mathrm{Del}_{\mathcal{M}}\, P$ has homeomorphic underlying space to $\mathcal{M}$.

Topology of RDT

TBP Theorem [Edelsbrunner-Shah 94] : If $\mathrm{Vor}\, P$ has the topological ball property w.r.t. $\mathcal{M}$, then $\mathrm{Del}_{\mathcal{M}}\, P$ has homeomorphic underlying space to $\mathcal{M}$.

RDT Theorem [AB98, CDES01] : If P is 0.2-sample for a surface $\mathcal{M} \subset \mathbb{R}^3$, then $\mathrm{Vor}\, P$ satisfies TBP with respect to $\mathcal{M}$.

Difficulty I: No RDT Thm.

Negative result [Cheng-Dey-Ramos 05] : For a k-manifold $\mathcal{M} \subset \mathbb{R}^d$, $\mathrm{Del}_{\mathcal{M}}\, P$ may not be homeomorphic to $\mathcal{M}$ no matter how dense P is when $k > 2$ and $d > 3$.

- Due to this result, Witness complex [Carlsson-de Silva] may not be homeomorphic to $\mathcal{M}$ as noted in [Boissonnat-Oudot-Guibas 07].

Difficulty I: No RDT Thm.

Negative result [Cheng-Dey-Ramos 05] : For a k-manifold $\mathcal{M} \subset \mathbb{R}^d$, $\mathrm{Del}_{\mathcal{M}} P$ may not be homeomorphic to $\mathcal{M}$ no matter how dense P is when $k > 2$ and $d > 3$.

- Due to this result, Witness complex [Carlsson-de Silva] may not be homeomorphic to $\mathcal{M}$ as noted in [Boissonnat-Oudot-Guibas 07].

  - Problem caused by slivers

Difficulty I: No RDT Thm.

Negative result [Cheng-Dey-Ramos 05] : For a k-manifold $\mathcal{M} \subset \mathbb{R}^d$, $\mathrm{Del}_{\mathcal{M}}\, P$ may not be homeomorphic to $\mathcal{M}$ no matter how dense P is when $k > 2$ and $d > 3$.

- Due to this result, Witness complex [Carlsson-de Silva] may not be homeomorphic to $\mathcal{M}$ as noted in [Boissonnat-Oudot-Guibas 07].

- Problem caused by slivers

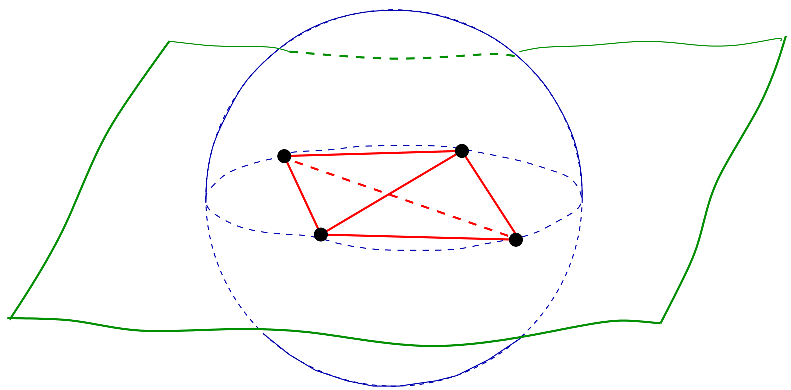

Difficulty I

- RDT theorem used the fact that Voronoi faces intersect $\mathcal{M}$ orthogonally.
- This is not true in high dimensions because of slivers whose dual faces may have large deviations from the normal space.

Simplex Shape

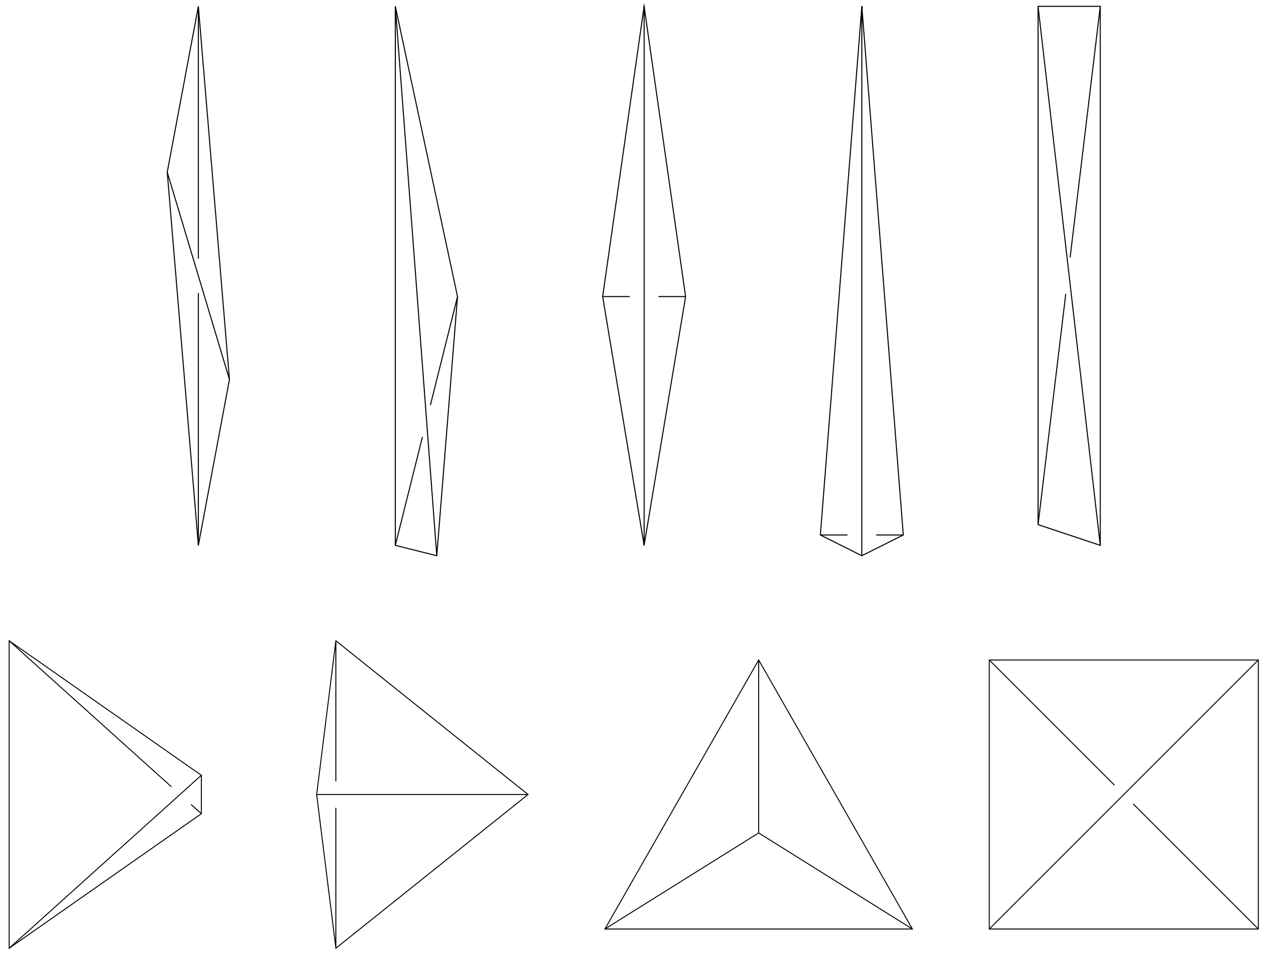

Simplex Shape

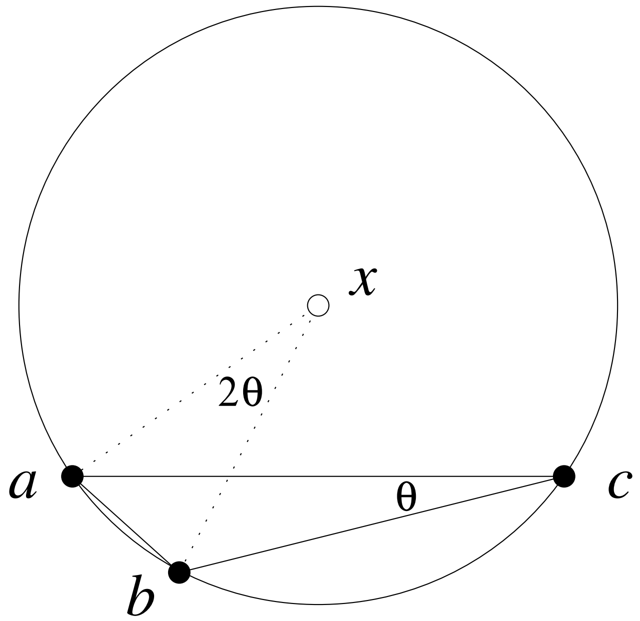

$$\sin\theta = \frac{L_\tau}{2R_\tau}$$

R	$\ldots$	circumradius
L_τ	$\ldots$	shortest edge length
R_τ	$\ldots$	circumradius
R_τ / L_τ	$\ldots$	circumradius-edge ratio

Sliver

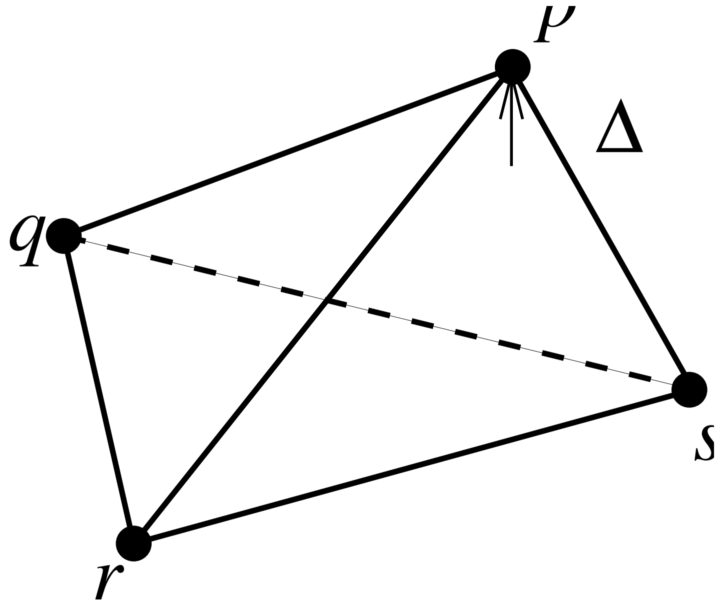

A j-simplex, $j > 1$, τ is a sliver if none of its subsimplices are sliver and

$$vol(\tau) \leq \sigma^j L_\tau^j$$

where L_τ is the shortest edge length, and σ is a parameter.

Difficulty I: Slivers in 3D

In 3-d, if a Delaunay triangle has a circumradius $O(\varepsilon f(p))$ then its normal and the normal of $\mathcal{M}$ at p form an angle $O(\varepsilon)$.

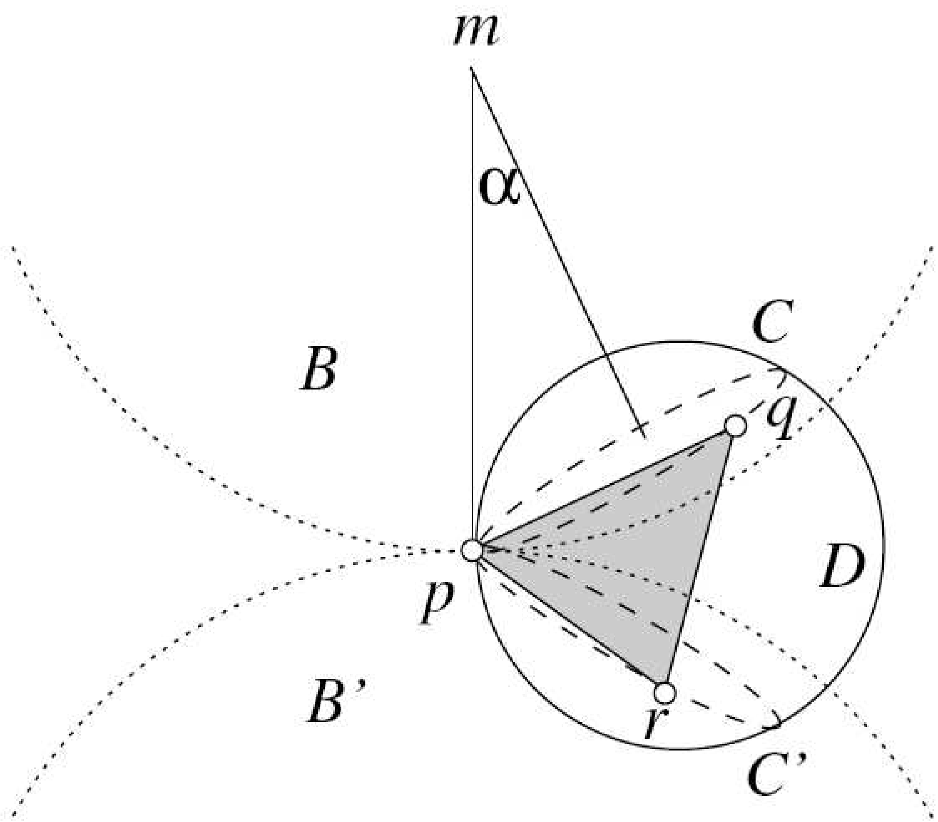

Difficulty I: Slivers' Normals

In 4-d, considering a 3-manifold, a Delaunay 3-simplex may have small circumradius, that is $O(\varepsilon f(p))$, but its normal may be very different from that of $\mathcal{M}$ at p. Slivers are the culprits:

$$p = (0, 0, \Delta, 0); q = (1, 0, 0, 0), r = (1, 1, 0, 0), s = (0, 1, 0, 0)$$

$$p = (0, 0, 0, \Delta); q = (1, 0, 0, 0), r = (1, 1, 0, 0), s = (0, 1, 0, 0)$$

Bad Normal

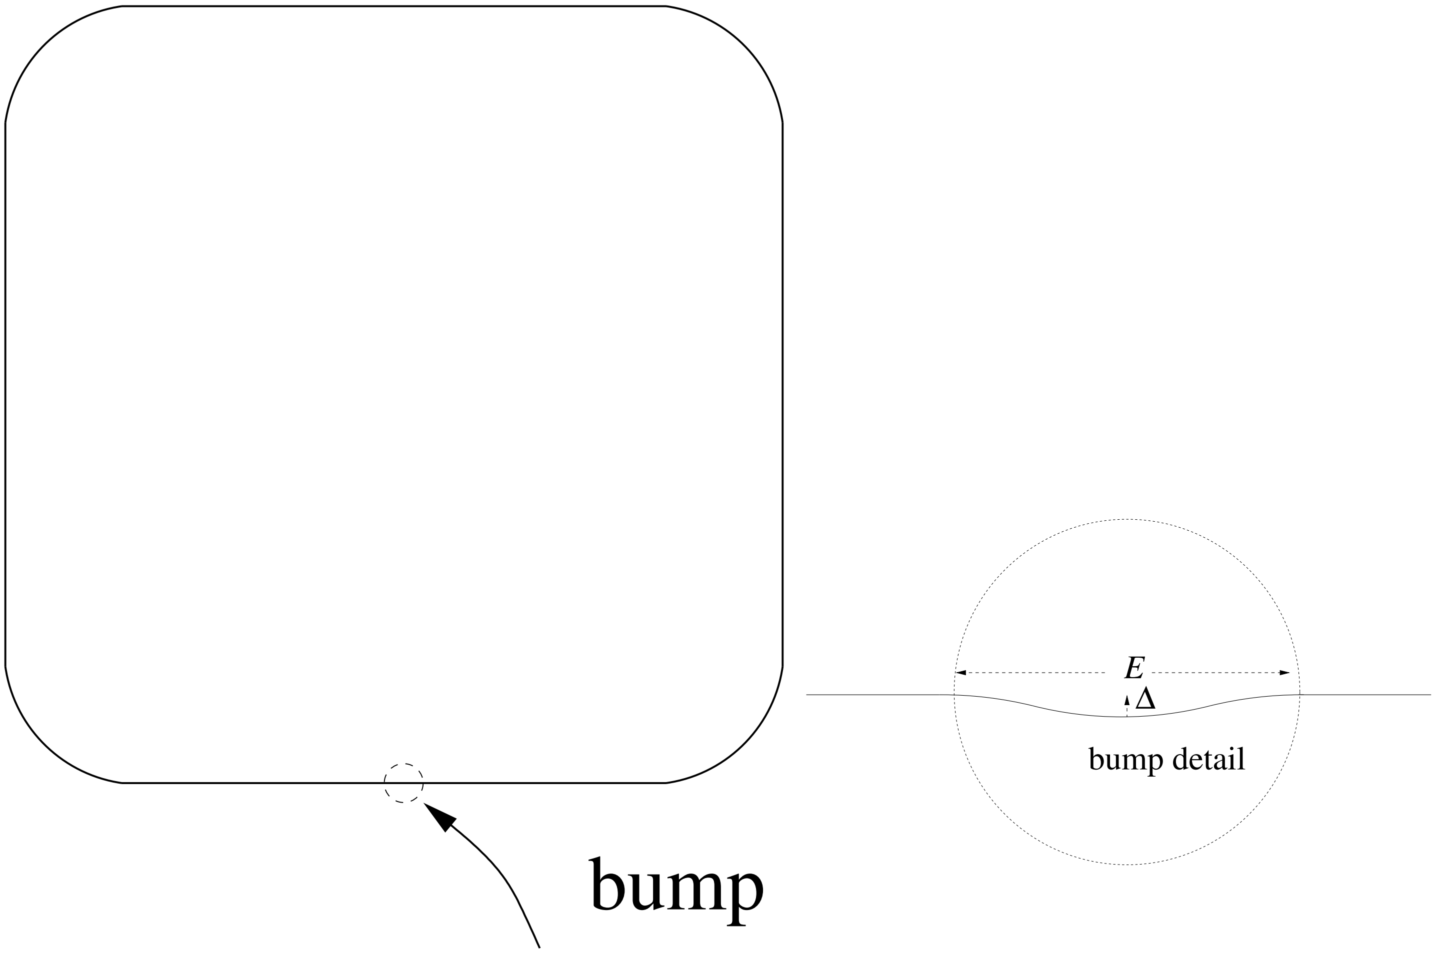

Difficulty II

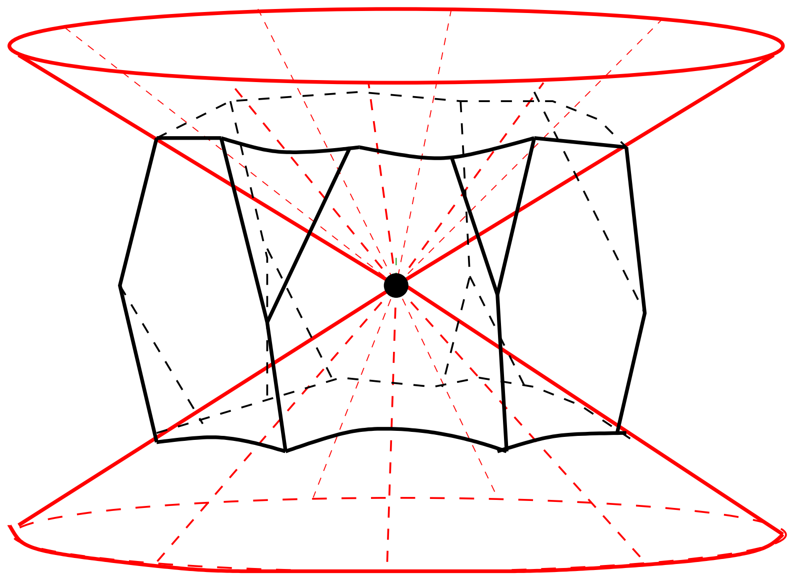

Difficulty II

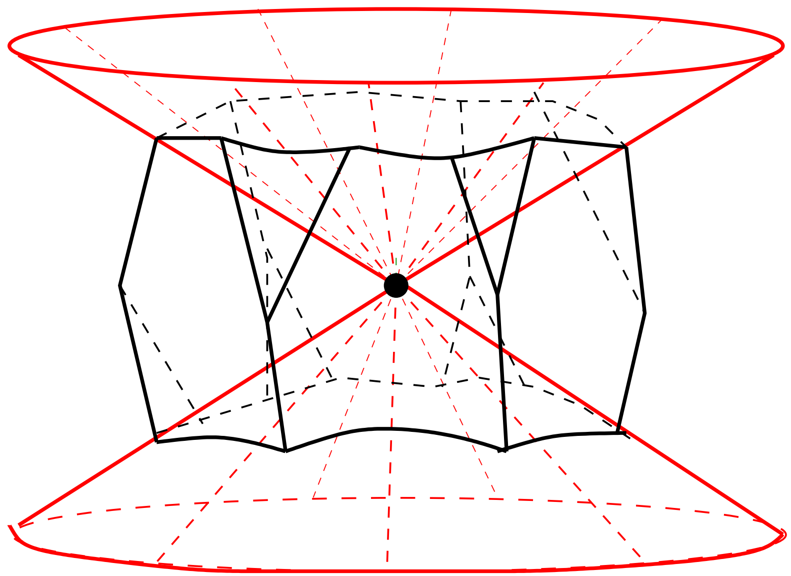

- It is not possible to identify precisely the restricted Delaunay triangulation because of slivers: there are Voronoi faces close to the surface but not intersecting it.

Solution[Cheng-Dey-Ramos 05]

Get rid of the slivers:

Follow sliver exudation approach of Cheng-Dey-Edelsbrunner-Facello-Teng 2000 in the context of meshing.

Weighted Voronoi/Delaunay

Weighted Points: (p, w_p).

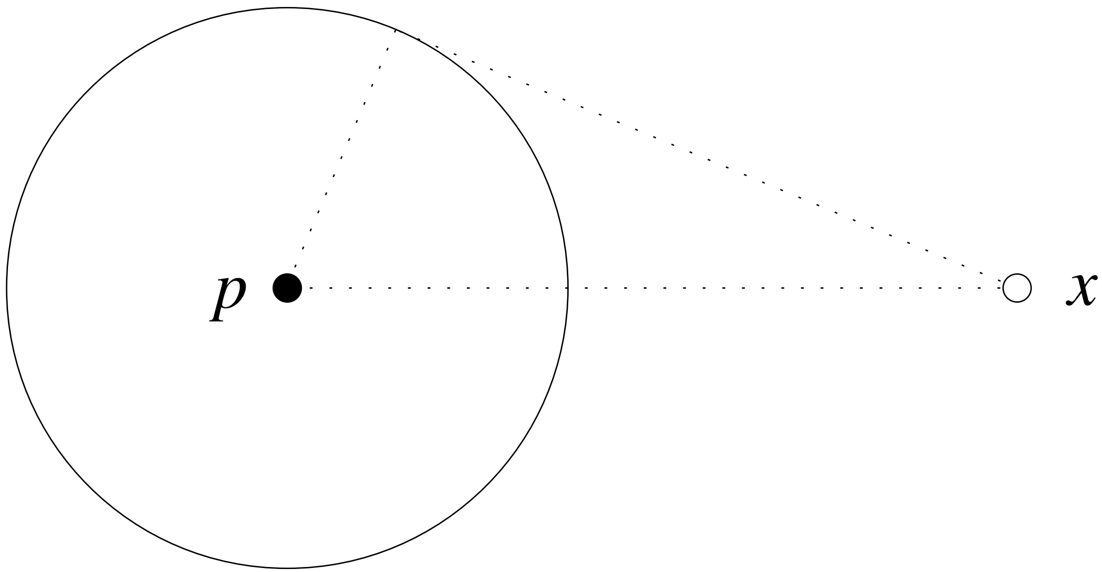

Weighted Distance: $\pi_{(p,w_p)}(x) = \|x - p\|^2 - w_p^2$.

Weighted Voronoi/Delaunay

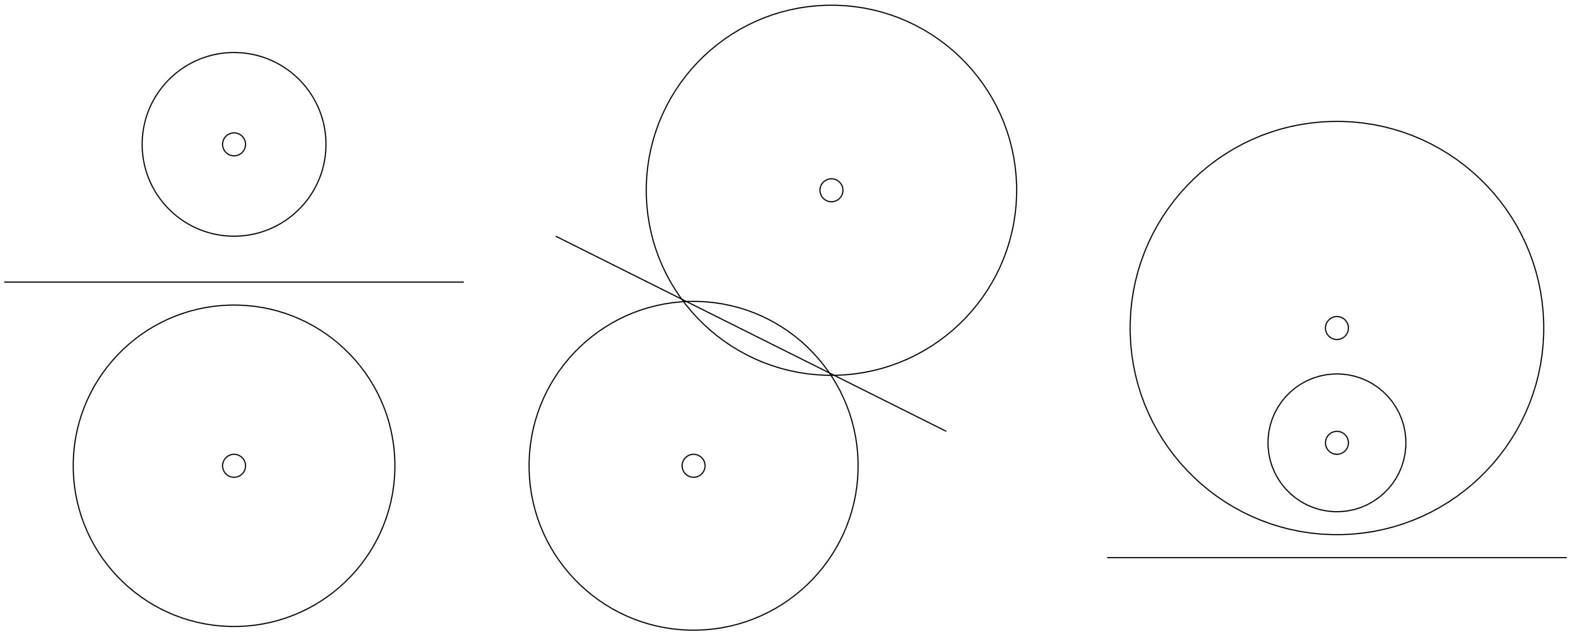

Bisectors of weighted points

Weight property $[\omega]$: For each $p \in P$, $w_p \leq \omega N(p)$, where $N(p)$ is the distance to closest neighbor. Limit $\omega \leq 1/4$.

Weighted Voronoi/Delaunay

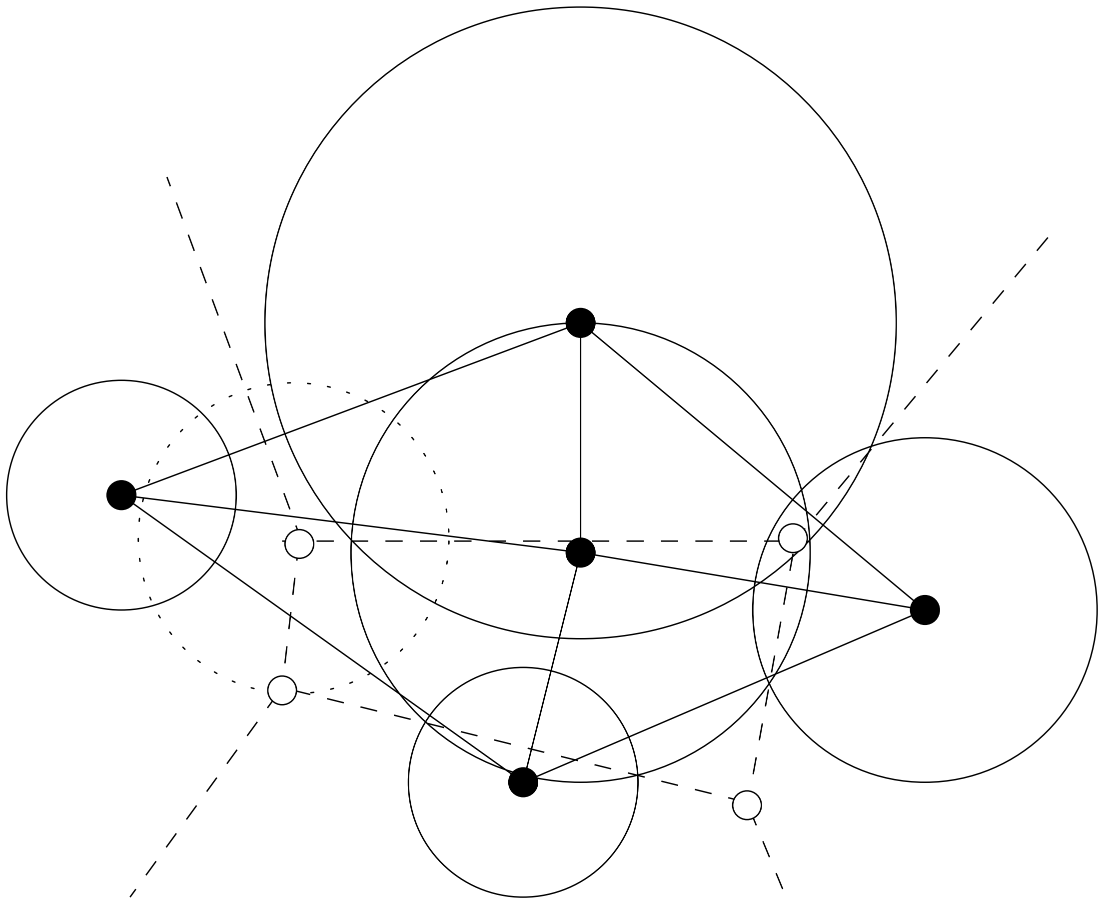

Weighted Delaunay/Voronoi complexes

Orthocenter Sensitivity

- An appropriate weight on each sample moves undesirable Voronoi faces away from the manifold

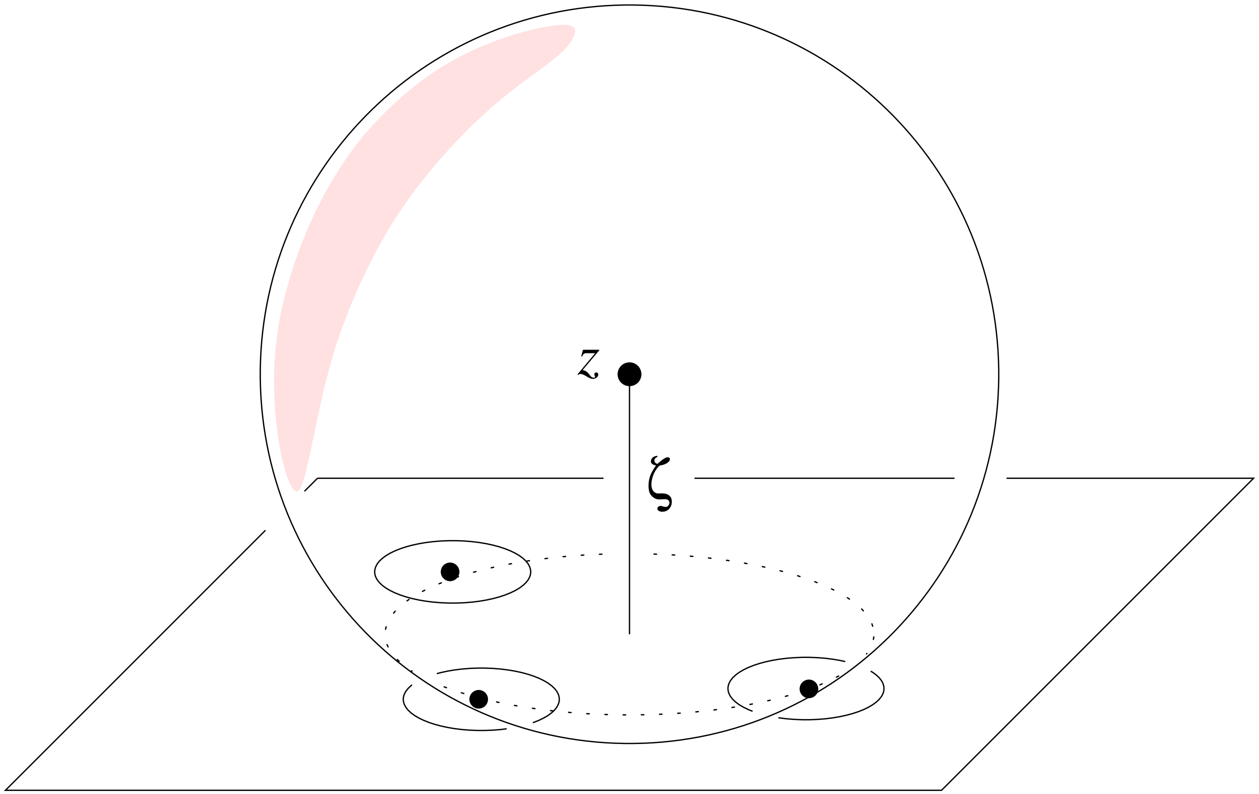

Sliver Exudation

- By a packing argument, the number of neighbors of p that can be connected to p in any weighted Delaunay triangulation with $[\omega]$ weight property is

$$\lambda = O(\nu^{2d})$$

  where $\nu = \varepsilon/\delta$.

- The number of possible weighted Delaunay $(\leq k)$-simplices in which p is involved is at most

$$N_p = O(\lambda^k) = O(\nu^{2kd})$$

Sliver Exudation

- The total w_p^2 length of “bad” intervals is at most:

$$N_p \cdot c' \sigma \varepsilon^2 f^2(p).$$

- The total w_p^2 length is $\omega^2 \delta^2 f^2(p)$. So if we choose

$$\sigma < \frac{\omega^2}{c' N_p \nu^2}$$

  then there is a radius w_p for p such that no cocone simplex is a sliver.

- For ε sufficiently small, if τ is a cocone $(k+1)$simplex, it must be a sliver.

Algorithm

- Construct $\mathrm{Vor}\ P$ and $\mathrm{Del}\ P$.
- Determine the dimension k of $\mathcal{M}$.
- “Pump up” the sample point weights to remove all j-slivers, $j = 3, \ldots, k+1$, from all point cocones.
- Extract all cocone simplices as the resulting output.

Correctness

The algorithm outputs $\mathrm{Del}_{\mathcal{M}}(\widehat{P})$:

(i) for σ sufficiently small, there is a weight assignment to the sample points so that no cocone j-simplex, $j \leq k+1$, is a sliver;

(ii) for ε sufficiently small, any cocone $(k+1)$-simplex must be a sliver.

Correctness

$\mathrm{Del}_{\mathcal{M}}(\widehat{P})$ is "close" to $\mathcal{M}$ and homeomorphic to it:

(i) the normal spaces of close points in $\mathcal{M}$ are close: if p, q are at distance $O(\varepsilon f(p))$, then their normal spaces form an angle $O(\varepsilon)$;

(ii) for any j-simplex with $j \leq k$, if its circumradius is $O(\varepsilon f(p))$ and neither τ nor any of the boundary simplices is a sliver, then the normal space of τ is close to the normal space of $\mathcal{M}$ at p;

(iii) each cell of $\mathrm{Vor}_{\mathcal{M}}\,\widehat{P}$ is a topological ball

(iii) implies that $\mathrm{Del}_{\mathcal{M}}\,\widehat{P}$ is homeomorphic to $\mathcal{M}$ (Edelsbrunner and Shah)

Summary

- For sufficiently dense samples, the algorithm outputs a mesh that is *faithful* to the original manifold.
- The running time – under (ε, δ)-sampling – is $O(n \log n)$ (constant is exponential with the dimension).
- The ε for which the algorithm works is quite small (more than exponentially small in the dimension).

Concluding remarks

- Various connections to Voronoi diagrams and geometry/topology estimations
- Further works on manifold/compact reconstructions [Niyogi-Smale-Weinberger 06, Chazal-Lieutier 06, Chazal-Cohen Steiner-Lieutire 06, Boissonnat-Oudot-Guibas 07]
- Practical algorithm under realistic assumptions ...???

Thank You