Good Triangulations

Jean-Daniel Boissonnat
DataShape, INRIA
http://www-sop.inria.fr/geometrica

Definition and existence of nets

Definition

Let Ω be a bounded subset of $\mathbb{R}^d$. A finite set of points P is called an $(\varepsilon, \bar{\eta})$-net of Ω iff

Density : $\forall x \in \Omega, \exists p \in P : \|x - p\| \leq \varepsilon$

Separation : $\forall p, q \in P : \|p - q\| \geq \bar{\eta}\,\varepsilon$

Lemma Ω admits an $(\varepsilon, 1)$-net.

Proof. While there exists a point $p \in \Omega$, $d(p, P) \geq \varepsilon$, insert p in P

Size of a net

Lemma The number of points of an $(\varepsilon, \bar{\eta})$-net is at most

$$n(\varepsilon, \bar{\eta}) \leq \frac{\mathrm{vol}_d(\Omega^{\frac{\eta}{2}})}{\mathrm{vol}_d(B(\frac{\eta}{2}))} = O\left(\frac{1}{\varepsilon^d}\right)$$

where the constant in the O depends on the geometry of Ω and on $\bar{\eta}^d$.

Proof. Consider the balls $B(p, \frac{\eta}{2})$ of radius $\frac{\eta}{2}$ that are centered at the points $p \in$ P. These balls are disjoint by definition of an $(\varepsilon, \bar{\eta})$-sample and they are all contained in $\Omega^{\frac{\eta}{2}}$

Delaunay complex of a net

lemma Let Ω be a bounded subset of $\mathbb{R}^d$, P an $(\varepsilon, \bar{\eta})$-net of Ω, and assume that d and $\bar{\eta}$ are positive constants. The restriction of the Delaunay triangulation of P to Ω has linear size $O(n)$ where $n = |\mathsf{P}| = O(\frac{1}{\varepsilon^d})$

Proof. 1. First bound the number of neighbours of p is $n_p = O(1)$ using a volume argument

2. Bound the number of simplices incident on a vertex is at most

$$\sum_{i=1}^{d+1} \binom{n_p}{i} \leq \sum_{i=0}^{n_p} \binom{n_p}{i} = 2^{n_p}.$$

3. For the construction, use a grid G_ε of resolution ε and compute, for each $p \in P$, the subset $N(p) \subset P$ of points that lie at distance at most 2ε from the cell that contains p. We have

$$|N(p)| = O(1) \quad \text{and} \quad \mathrm{star}(p, \mathrm{Del}_{|\Omega}(\mathsf{P})) = \mathrm{star}(p, \mathrm{Del}_{|\Omega}(N(p)))$$

Nets in discrete metric spaces

We are only given the distance of interpoint distances (not the locations of the points)

Lemma Let W be a finite set of points such that the distance of any point $q \in W$ to $W \setminus \{q\}$ is at most ε and let $\lambda \geq \varepsilon$. One can extract from W a subsample L that is a $(\lambda, 1)$-net of W.

Farthest point insertion

Input: the distance matrix of a finite point set W and either a positive constant λ (Case 1) or an integer k (Case 2)

1. $L := \emptyset$
2. $L(w) := p_\infty$ for all $w \in W$
3. $\lambda^* := \max_{w \in W} \|w - L(w)\|$
4. $w^* :=$ a point $p \in W$ such that $\|p - L(p)\| = \lambda^*$
5. **while** either $\lambda^* > \lambda$ (Case 1) or $|L| < k$ (Case 2)
   5.1 add w^* to L
   5.2 **for** each point w of W such that $\|w - w^*\| < \|w - L(w)\|$ **do**
      5.2.1 $L(w)) := w^*$
      5.2.2 update w^*
6. **Output :** $L \subseteq W$, a $(\lambda, 1)$-net of W (Case 1), an approximate solution to the k-centers problem (Case 2)

Analysis of the algorithm

For any $i > 0$, $L_i = \{p_1, ..., p_i\}$ and $\lambda_i = d(p_i, L_{i-1})$

Since L_i grows with i : $\quad j \geq i \ \Rightarrow \ \lambda_j \leq \lambda_i$

Lemma At each iteration $i > 0$, L_i is a $(\lambda_i, 1)$-net of W.

Proof

1. L_i is λ_i-dense in W ...
2. L_i is λ_i-separated: $\quad p_a p_b$ closest par in L_i, $\|p_a - p_b\| = \lambda_b \geq \lambda_i$

The k-centers problem

Problem : Select from W a subset L of k points so as to maximize the minimum pairwise distance between the points of L.

Lemma The farthest insertion algorithm (Case 2) provides a 2-approximation to the k-centers problem.

Proof

- $W \subset \cup_{i=1}^{k-1} B(l_i, \lambda_k)$

  $\Rightarrow$ Two points of L_{opt} lie in the same ball $B(l_i, \lambda_k), \quad i \leq k-1$

  $\Rightarrow \exists p, q \in L_{\text{opt}}$ s.t. $\|p-q\| \leq 2\lambda_k$

- The distance between any two points of L is at least λ_k.

Some optimality properties of Delaunay triangulations

Among all possible triangulations of $\mathcal{P}$, $\mathrm{Del}(\mathcal{P})$

1. (2d) maximizes the smallest angle [Lawson]

2. (2d) Linear interpolation of $\{(p_i, f(p_i))\}$ that minimizes [Rippa]

   $R(T) = \sum_i \int_{T_i} \left(\left(\frac{\partial \phi_i}{\partial x}\right)^2 + \left(\frac{\partial \phi_i}{\partial y}\right)^2 \right) dx\, dy$ (Dirichlet energy)

   $\phi_i =$ linear interpolation of the $f(p_j)$ over triangle $T_i \in T$

3. minimizes the radius of the maximal smallest ball enclosing a simplex) [Rajan]

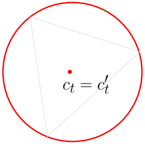

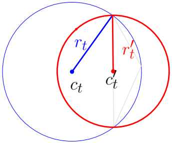

Some optimality properties of Delaunay triangulations

Among all possible triangulations of $\mathcal{P}$, $\mathrm{Del}(\mathcal{P})$

1. (2d) maximizes the smallest angle [Lawson]

2. (2d) Linear interpolation of $\{(p_i, f(p_i))\}$ that minimizes [Rippa]

   $$R(T) = \sum_i \int_{T_i} \left(\left(\frac{\partial \phi_i}{\partial x}\right)^2 + \left(\frac{\partial \phi_i}{\partial y}\right)^2 \right) dx\, dy \qquad \text{(Dirichlet energy)}$$

   $\phi_i =$ linear interpolation of the $f(p_j)$ over triangle $T_i \in T$

3. minimizes the radius of the maximal smallest ball enclosing a simplex) [Rajan]

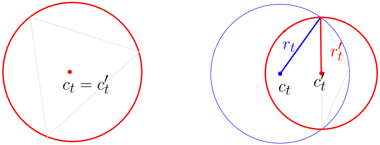

Some optimality properties of Delaunay triangulations

Among all possible triangulations of $\mathcal{P}$, $\mathrm{Del}(\mathcal{P})$

1. (2d) maximizes the smallest angle [Lawson]

2. (2d) Linear interpolation of $\{(p_i, f(p_i))\}$ that minimizes [Rippa]

   $R(T) = \sum_i \int_{T_i} \left(\left(\frac{\partial \phi_i}{\partial x}\right)^2 + \left(\frac{\partial \phi_i}{\partial y}\right)^2 \right) dx\, dy$ (Dirichlet energy)

   $\phi_i =$ linear interpolation of the $f(p_j)$ over triangle $T_i \in T$

3. minimizes the radius of the maximal smallest ball enclosing a simplex) [Rajan]

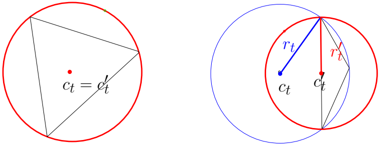

Optimizing the angular vector ($d = 2$)

Angular vector of a triangulation $T(\mathcal{P})$

$$\text{ang}\,(T(\mathcal{P})) = (\alpha_1, \ldots, \alpha_{3t}), \qquad \alpha_1 \leq \ldots \leq \alpha_{3t}$$

Optimality Any triangulation of a given point set $\mathcal{P}$ whose angular vector is maximal (for the lexicographic order) is a Delaunay triangulation of $\mathcal{P}$

Good for matrix conditioning in FE methods

Local characterization of Delaunay complexes

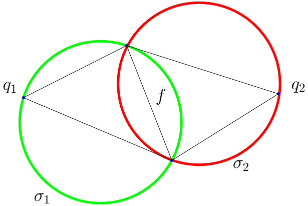

Pair of regular simplices

$$\sigma_2(q_1) \geq 0 \quad \text{and} \quad \sigma_1(q_2) \geq 0$$

$\Leftrightarrow \hat{c}_1 \in h^+_{\sigma_2}$ and $\hat{c}_2 \in h^+_{\sigma_1}$

Theorem A triangulation $T(P)$ such that all pairs of simplexes are regular is a Delaunay triangulation $\mathrm{Del}(P)$

Proof The PL function whose graph G is obtained by lifting the triangles is locally convex and has a convex support

$$\Rightarrow \quad G = \mathrm{conv}^-(\hat{Q}) \quad \Rightarrow \quad T(Q) = \mathrm{Del}(Q)$$

Local characterization of Delaunay complexes

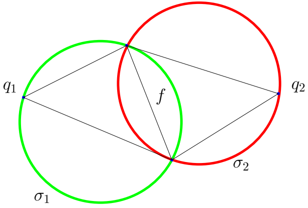

Pair of regular simplices

$$\sigma_2(q_1) \geq 0 \quad \text{and} \quad \sigma_1(q_2) \geq 0$$

$\Leftrightarrow \hat{c}_1 \in h^+_{\sigma_2}$ and $\hat{c}_2 \in h^+_{\sigma_1}$

Theorem A triangulation $T(P)$ such that all pairs of simplexes are regular is a Delaunay triangulation $\mathrm{Del}(P)$

Proof The PL function whose graph G is obtained by lifting the triangles is locally convex and has a convex support

$$\Rightarrow \quad G = \mathrm{conv}^-(\hat{Q}) \quad \Rightarrow \quad T(Q) = \mathrm{Del}(Q)$$

Local characterization of Delaunay complexes

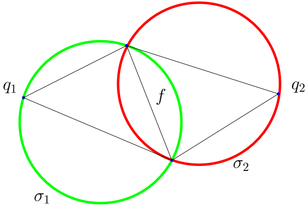

Pair of regular simplices

$$\sigma_2(q_1) \geq 0 \quad \text{and} \quad \sigma_1(q_2) \geq 0$$

$\Leftrightarrow \hat{c}_1 \in h^+_{\sigma_2}$ and $\hat{c}_2 \in h^+_{\sigma_1}$

Theorem A triangulation $T(P)$ such that all pairs of simplexes are regular is a Delaunay triangulation $\mathrm{Del}(P)$

Proof The PL function whose graph G is obtained by lifting the triangles is locally convex and has a convex support

$$\Rightarrow \quad G = \mathrm{conv}^-(\hat{Q}) \quad \Rightarrow \quad T(Q) = \mathrm{Del}(Q)$$

Lawson's proof using flips

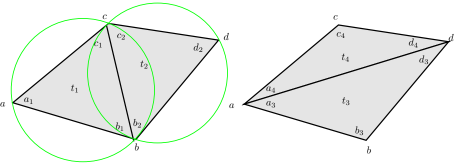

While $\exists$ a non regular pair (t_3, t_4)

/* $t_3 \cup t_4$ is convex */

replace (t_3, t_4) by (t_1, t_2)

Regularize $\Leftrightarrow$ improve $\mathrm{ang}\,(T(\mathcal{P}))$

$\mathrm{ang}\,(t_1, t_2) \geq \mathrm{ang}\,(t_3, t_4)$

$$a_1 = a_3 + a_4,\ d_2 = d_3 + d_4,$$
$$c_1 \geq d_3, \quad b_1 \geq d_4, \quad b_2 \geq a_4, \quad c_2 \geq a_3$$

- The algorithm terminates since the number of triangulations of $\mathcal{P}$ is finite and $\mathrm{ang}(T(\mathcal{P}))$ cannot decrease
- The obtained triangulation is a Delaunay triangulation of $\mathcal{P}$ since all its edges are regular

Lawson's proof using flips

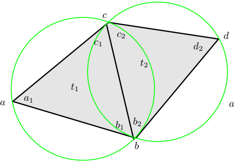

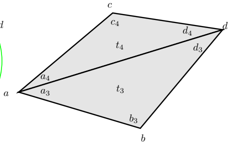

While $\exists$ a non regular pair (t_3, t_4)

/* $t_3 \cup t_4$ is convex */

replace (t_3, t_4) by (t_1, t_2)

Regularize $\Leftrightarrow$ improve $\text{ang}\,(T(\mathcal{P}))$

$\text{ang}\,(t_1, t_2) \geq \text{ang}\,(t_3, t_4)$

$$a_1 = a_3 + a_4,\ d_2 = d_3 + d_4,$$
$$c_1 \geq d_3, \quad b_1 \geq d_4, \quad b_2 \geq a_4, \quad c_2 \geq a_3$$

- The algorithm terminates since the number of triangulations of $\mathcal{P}$ is finite and $\text{ang}(T(\mathcal{P}))$ cannot decrease
- The obtained triangulation is a Delaunay triangulation of $\mathcal{P}$ since all its edges are regular

Lawson's proof using flips

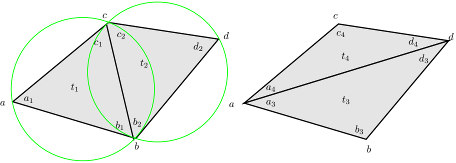

While $\exists$ a non regular pair (t_3, t_4)

/* $t_3 \cup t_4$ is convex */

replace (t_3, t_4) by (t_1, t_2)

Regularize $\Leftrightarrow$ improve $\text{ang}\,(T(\mathcal{P}))$

$\text{ang}\,(t_1, t_2) \geq \text{ang}\,(t_3, t_4)$

$$a_1 = a_3 + a_4,\ d_2 = d_3 + d_4,$$
$$c_1 \geq d_3, \quad b_1 \geq d_4, \quad b_2 \geq a_4, \quad c_2 \geq a_3$$

- The algorithm terminates since the number of triangulations of $\mathcal{P}$ is finite and $\text{ang}(T(\mathcal{P}))$ cannot decrease
- The obtained triangulation is a Delaunay triangulation of $\mathcal{P}$ since all its edges are regular

Flat simplices may exist in higher dimensional DT

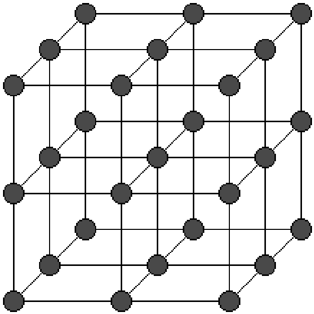

- Each square face can be circumscribed by an empty sphere
- This remains true if the grid points are slightly perturbed therefore creating flat tetrahedra

The long quest for thick triangulations

Differential Topology [Cairns], [Whitehead], [Whitney], [Munkres]

Differential Geometry [Cheeger et al.]

Geometric Function Theory [Peltonen], [Saucan]

Simplex quality

Altitudes

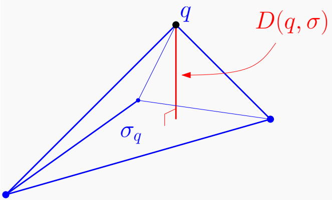

If σ_q, the face opposite q in σ is protected, The *altitude* of q in σ is

$$D(q,\sigma) = d(q, \mathrm{aff}(\sigma_q)),$$

where σ_q is the face opposite q.

Definition (Thickness [Cairns, Whitney, Whitehead et al.])

The *thickness* of a j-simplex σ with diameter $\Delta(\sigma)$ is

$$\Theta(\sigma) = \begin{cases} 1 & \text{if } j = 0 \\ \min_{p\in\sigma} \frac{D(p,\sigma)}{j\Delta(\sigma)} & \text{otherwise.} \end{cases}$$

Simplex quality

Altitudes

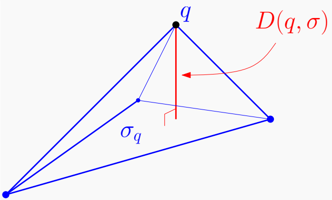

If σ_q, the face opposite q in σ is protected, The *altitude* of q in σ is

$$D(q,\sigma) = d(q, \mathrm{aff}(\sigma_q)),$$

where σ_q is the face opposite q.

Definition (Thickness [Cairns, Whitney, Whitehead et al.])

The *thickness* of a j-simplex σ with diameter $\Delta(\sigma)$ is

$$\Theta(\sigma) = \begin{cases} 1 & \text{if } j = 0 \\ \min_{p\in\sigma} \frac{D(p,\sigma)}{j\Delta(\sigma)} & \text{otherwise.} \end{cases}$$

Protection

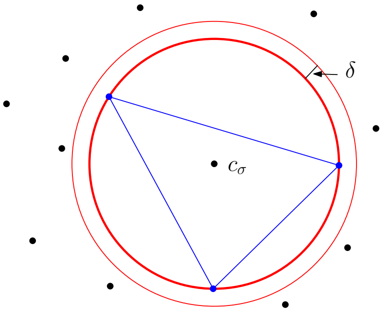

δ-protection We say that a Delaunay simplex $\sigma \subset L$ is δ-protected if

$$\|c_\sigma - q\| > \|c_\sigma - p\| + \delta \qquad \forall p \in \sigma \ \text{ and } \ \forall q \in L \setminus \sigma.$$

Protection implies separation and thickness

Let P be a $(\varepsilon, \bar{\eta})$-net, i.e.

- $\forall x \in \Omega, \quad d(x,P) \leq \varepsilon$
- $\forall p,q \in P, \quad \|p-q\| \geq \bar{\eta}\varepsilon$

if all d-simplices of $\mathrm{Del}(P)$ are $\bar{\delta}\varepsilon$-protected, then

- **Separation** of P : $\bar{\eta} \geq \bar{\delta}$
- **Thickness :** $\forall \sigma \in \mathrm{Del}(P), \quad \Theta(\sigma) \geq \frac{\bar{\delta}^2}{8d}$

Given: A set of (bad) events $A_1, ..., A_N$,
each happens with $\text{proba}(A_i) \leq p < 1$

Question : what is the probability that none of the events occur?

The case of independent events

$$\text{proba}(\neg A_1 \wedge ... \wedge \neg A_N) \geq (1-p)^N > 0$$

What if we allow a limited amount of dependency among the events?

Given: A set of (bad) events $A_1, ..., A_N$,
each happens with $\text{proba}(A_i) \leq p < 1$

Question : what is the probability that none of the events occur?

The case of independent events

$$\text{proba}(\neg A_1 \wedge ... \wedge \neg A_N) \geq (1-p)^N > 0$$

What if we allow a limited amount of dependency among the events?

Given: A set of (bad) events $A_1, ..., A_N$,
each happens with $\text{proba}(A_i) \leq p < 1$

Question : what is the probability that none of the events occur?

The case of independent events

$$\text{proba}(\neg A_1 \wedge ... \wedge \neg A_N) \geq (1-p)^N > 0$$

What if we allow a limited amount of dependency among the events?

LLL : symmetric version [Lovász & Erdös 1975]

Under the assumptions

1. $\text{proba}(A_i) \leq p < 1$
2. A_i depends of $\leq \Gamma$ other events A_j
3. $\text{proba}(A_i) \leq \frac{1}{e\,(\Gamma+1)}$ $\qquad e = 2.718...$

then

$$\text{proba}(\neg A_1 \wedge ... \wedge \neg A_N) > 0$$

Moser and Tardos' constructive proof of the LLL [2010]

$\mathcal{P}$ a finite set of mutually independent random variables

$\mathcal{A}$ a finite set of events that are determined by the values of $S \subseteq \mathcal{P}$

Two events are independent iff they share no variable

Algorithm

for all $P \in \mathcal{P}$ **do**
    $v_P \leftarrow$ a random evaluation of P;

while $\exists A \in \mathcal{A}$: A happens when $(P = v_P, P \in \mathcal{P})$ **do**

    pick an arbitrary happening event $A \in \mathcal{A}$;

    for all $P \in \mathrm{variables}(A)$ **do**
        $v_P \leftarrow$ a new random evaluation of P;

return $(v_P)_{P \in \mathcal{P}}$;

Moser and Tardos' constructive proof of the LLL [2010]

$\mathcal{P}$ a finite set of mutually independent random variables

$\mathcal{A}$ a finite set of events that are determined by the values of $S \subseteq \mathcal{P}$

Two events are independent iff they share no variable

Algorithm

for all $P \in \mathcal{P}$ **do**
  $v_P \leftarrow$ a random evaluation of P;

while $\exists A \in \mathcal{A}$: A happens when $(P = v_P, P \in \mathcal{P})$ **do**

  pick an arbitrary happening event $A \in \mathcal{A}$;

  for all $P \in \mathrm{variables}(A)$ **do**
    $v_P \leftarrow$ a new random evaluation of P;

return $(v_P)_{P \in \mathcal{P}}$;

Moser and Tardos' theorem

if

1. $\text{proba}(A_i) \leq p < 1$
2. A_i depends of $\leq \Gamma$ other events A_j
3. $\text{proba}(A_i) \leq \frac{1}{e\,(\Gamma+1)}$ $\qquad e = 2.718...$

then $\exists$ an assignment of values to the variables $\mathcal{P}$ such that no event in $\mathcal{A}$ happens

The randomized algorithm resamples an event $A \in \mathcal{A}$ at most $\boxed{\frac{1}{\Gamma}}$ expected times before it finds such an evaluation

The expected total number of resampling steps is at most $\boxed{\frac{N}{\Gamma}}$

Home work

- Read the proof of Moser & Tardos (or Spencer's nice note)
- Learn about the parallel and the derandomized versions
- Listen to a talk by Aravind Srinivasan on further extensions

  `https://video.ias.edu/csdm/2014/0407-AravindSrinivasan`

Protecting Delaunay simplices via perturbation

Picking regions : pick p' in $B(p, \rho)$ Hyp. $\rho \leq \frac{\eta}{4} \;\; (\leq \frac{\varepsilon}{2})$

Sampling parameters of a perturbed point set

If P is an $(\varepsilon, \bar{\eta})$-net, P′ is an $(\varepsilon', \bar{\eta}')$-net, where

$$\varepsilon' = \varepsilon(1 + \bar{\rho}) \;\text{ and }\; \bar{\eta}' = \frac{\bar{\eta} - 2\bar{\rho}}{1 + \bar{\rho}} \geq \frac{\bar{\eta}}{3}$$

Notation : $\bar{x} = \frac{x}{\varepsilon}$

Protecting Delaunay simplices via perturbation

Picking regions : pick p' in $B(p, \rho)$ Hyp. $\rho \leq \frac{\eta}{4} \ (\leq \frac{\varepsilon}{2})$

Sampling parameters of a perturbed point set

If P is an $(\varepsilon, \bar{\eta})$-net, P′ is an $(\varepsilon', \bar{\eta}')$-net, where

$$\varepsilon' = \varepsilon(1+\bar{\rho}) \text{ and } \bar{\eta}' = \frac{\bar{\eta} - 2\bar{\rho}}{1+\bar{\rho}} \geq \frac{\bar{\eta}}{3}$$

Notation : $\bar{x} = \frac{x}{\varepsilon}$

The LLL framework

Random variables : P' a set of random points $\{p',\ p' \in B(p,\rho), p \in \mathsf{P}\}$

Event: $\exists \phi' = (\sigma', p')$ (Bad configuration)

σ' is a d simplex with $R_{\sigma'} \leq \varepsilon + \rho$

$p' \in Z_\delta(\sigma') \qquad Z_\delta(\sigma') = B(c_{\sigma'}, R_{\sigma'} + \delta) \setminus B(c_{\sigma'}, R_{\sigma'})$

Algorithm

Input: P, ρ, δ

while an event $\phi' = (\sigma', p')$ occurs **do**

resample the points of ϕ'

update $\mathrm{Del}(\mathsf{P}')$

Output: P' and $\mathrm{Del}(\mathsf{P}')$

The LLL framework

Random variables : P' a set of random points $\{p',\ p' \in B(p,\rho), p \in \mathsf{P}\}$

Event: $\exists \phi' = (\sigma', p')$ (Bad configuration)

σ' is a d simplex with $R_{\sigma'} \leq \varepsilon + \rho$

$p' \in Z_\delta(\sigma') \qquad Z_\delta(\sigma') = B(c_{\sigma'}, R_{\sigma'} + \delta) \setminus B(c_{\sigma'}, R_{\sigma'})$

Algorithm

Input: P, ρ, δ

while an event $\phi' = (\sigma', p')$ occurs **do**

resample the points of ϕ'

update $\mathrm{Del}(\mathsf{P}')$

Output: P' and $\mathrm{Del}(\mathsf{P}')$

Bounding Γ

Lemma : An event is independent of all but at most Γ other bad events where Γ depends on $\bar{\eta}$, $\bar{\rho}$, $\bar{\delta}$ and d

Proof :

- Let $\phi' = (\sigma', p')$ be a bad configuration.

$$\forall p' \in \phi', \quad \|p' - c_{\sigma'}\| \le R_{\sigma'} + \delta = R = \varepsilon + \rho + \delta = \varepsilon\,(1 + \bar{\rho} + \bar{\delta})$$

- the number of events that may not be independent from an event (σ', p') is at most the number of subsets of $(d+1)$ points in $B(c_{\sigma'}, 3R)$.

- Since P′ is η'-sparse,

$$\Gamma = \left(\frac{3R + \frac{\eta'}{2}}{\frac{\eta'}{2}}\right)^{d(d+1)} = \left(1 + 6\,\frac{(1+\bar{\rho}+\bar{\delta})\,(1+\bar{\rho})}{\bar{\eta} - 2\bar{\rho}}\right)^{d(d+1)}$$

Bounding Γ

Lemma : An event is independent of all but at most Γ other bad events where Γ depends on $\bar{\eta}$, $\bar{\rho}$, $\bar{\delta}$ and d

Proof :

- Let $\phi' = (\sigma', p')$ be a bad configuration.

$$\forall p' \in \phi', \quad \|p' - c_{\sigma'}\| \le R_{\sigma'} + \delta = R = \varepsilon + \rho + \delta = \varepsilon\,(1 + \bar{\rho} + \bar{\delta})$$

- the number of events that may not be independent from an event (σ', p') is at most the number of subsets of $(d+1)$ points in $B(c_{\sigma'}, 3R)$.

- Since P′ is η'-sparse,

$$\Gamma = \left(\frac{3R + \frac{\eta'}{2}}{\frac{\eta'}{2}}\right)^{d(d+1)} = \left(1 + 6\,\frac{(1+\bar{\rho}+\bar{\delta})\,(1+\bar{\rho})}{\bar{\eta} - 2\bar{\rho}}\right)^{d(d+1)}$$

Bounding Γ

Lemma : An event is independent of all but at most Γ other bad events where Γ depends on $\bar{\eta}$, $\bar{\rho}$, $\bar{\delta}$ and d

Proof :

- Let $\phi' = (\sigma', p')$ be a bad configuration.

$$\forall p' \in \phi', \quad \|p' - c_{\sigma'}\| \leq R_{\sigma'} + \delta = R = \varepsilon + \rho + \delta = \varepsilon\,(1 + \bar{\rho} + \bar{\delta})$$

- the number of events that may not be independent from an event (σ', p') is at most the number of subsets of $(d+1)$ points in $B(c_{\sigma'}, 3R)$.

- Since P′ is η'-sparse,

$$\Gamma = \left(\frac{3R + \frac{\eta'}{2}}{\frac{\eta'}{2}} \right)^{d(d+1)} = \left(1 + 6\,\frac{(1 + \bar{\rho} + \bar{\delta})\,(1 + \bar{\rho})}{\bar{\eta} - 2\bar{\rho}} \right)^{d(d+1)}$$

Bounding Γ

Lemma : An event is independent of all but at most Γ other bad events where Γ depends on $\bar{\eta}$, $\bar{\rho}$, $\bar{\delta}$ and d

Proof :

- Let $\phi' = (\sigma', p')$ be a bad configuration.

$$\forall p' \in \phi', \quad \|p' - c_{\sigma'}\| \leq R_{\sigma'} + \delta = R = \varepsilon + \rho + \delta = \varepsilon\,(1 + \bar{\rho} + \bar{\delta})$$

- the number of events that may not be independent from an event (σ', p') is at most the number of subsets of $(d+1)$ points in $B(c_{\sigma'}, 3R)$.

- Since P′ is η'-sparse,

$$\Gamma = \left(\frac{3R + \frac{\eta'}{2}}{\frac{\eta'}{2}}\right)^{d(d+1)} = \left(1 + 6\,\frac{(1 + \bar{\rho} + \bar{\delta})\,(1 + \bar{\rho})}{\bar{\eta} - 2\bar{\rho}}\right)^{d(d+1)}$$

Bounding $\mathrm{proba}(\sigma, p)$ be a bad configuration

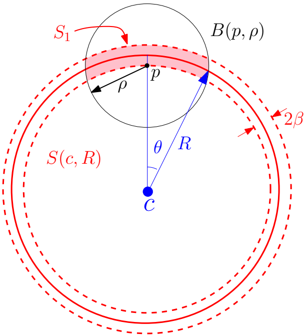

$S(c, R)$ a hypersphere of $\mathbb{R}^d$

$T_\delta = B(c, R + \delta) \setminus B(c, R)$

B_ρ any d-ball of radius $\rho < R$

$\mathrm{vol}_d(T_\delta \cap B_\rho) \leq U_{d-1}\,(R\theta)^{d-1}\,\delta,$

$\frac{2}{\pi}\theta \leq \sin\theta \leq \frac{\rho}{R} \qquad (\theta < \frac{\pi}{2} \Leftarrow\ \rho < R)$

$\Rightarrow \quad R\theta \leq \frac{\pi}{2}\rho$

$\mathrm{proba}(p' \in Z_\delta(\sigma')) \leq \varpi = \frac{U_{d-1}}{U_d}\,\frac{2}{\pi}\,\frac{\delta}{\rho} \leq \frac{C}{\sqrt{d}}\,\frac{\delta}{\rho}$

Bounding $\mathrm{proba}(\sigma, p)$ be a bad configuration

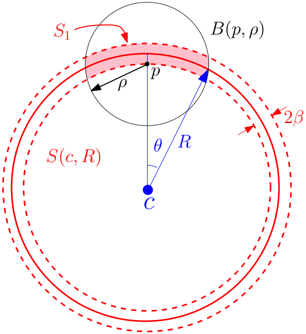

$S(c,R)$ a hypersphere of $\mathbb{R}^d$

$T_\delta = B(c, R+\delta) \setminus B(c,R)$

B_ρ any d-ball of radius $\rho < R$

$\mathrm{vol}_d(T_\delta \cap B_\rho) \leq U_{d-1}\,(R\theta)^{d-1}\,\delta,$

$\frac{2}{\pi}\theta \leq \sin\theta \leq \frac{\rho}{R} \qquad (\theta < \frac{\pi}{2} \Leftarrow\ \rho < R)$

$\Rightarrow\ \ R\theta \leq \frac{\pi}{2}\rho$

$$\mathrm{proba}(p' \in Z_\delta(\sigma')) \leq \varpi = \frac{U_{d-1}}{U_d}\,\frac{2}{\pi}\,\frac{\delta}{\rho} \leq \frac{C}{\sqrt{d}}\,\frac{\delta}{\rho}$$

Main result

Under condition

$$\frac{2^{d+1}e}{\pi}\,(\Gamma+1)\,\delta \leq \rho \leq \frac{\eta}{4}$$

the algorithm terminates.

Guarantees on the output

- $d_H(P, P') \leq \rho$
- the d-simplices of $\mathrm{Del}(P')$ are δ-protected
- and therefore have a positive thickness

Main result

Under condition

$$\frac{2^{d+1}e}{\pi}(\Gamma+1)\,\delta \leq \rho \leq \frac{\eta}{4}$$

the algorithm terminates.

Guarantees on the output

- $d_H(P, P') \leq \rho$
- the d-simplices of $\mathrm{Del}(P')$ are δ-protected
- and therefore have a positive thickness

Bound on the number of events

$\Sigma(p')$: number of d-simplices that can possibly make a bad configuration with $p' \in P'$ for some perturbed set P'

$R = \varepsilon + \rho + \delta$

$$\begin{aligned}\sum_{p' \in P'} \Sigma(p') &\leq n \times |P' \cap B(p', 2R)|^{d+1} \\ &\leq n \left(\frac{2\left(1 + \bar{\rho} + \bar{\delta} + \frac{\bar{\eta}'}{2}\right)}{\frac{\bar{\eta}'}{2}} \right)^{d(d+1)} \\ &= C' n\end{aligned}$$

Complexity of the algorithm

- The number of resamplings executed by the algorithm is at most

$$\frac{Cn}{\Gamma} \leq C'' n$$

  where C'' depends on $\bar{\eta}$, $\bar{\rho}$, $\bar{\delta}$ and (exponentially) d

- Each resampling consists in perturbing $O(1)$ points
- Updating the Delaunay triangulation after each resampling takes $O(1)$ time
- The expected complexity is linear in the number of points