

A Voronoi diagram is shown on a white background with a light yellow border. It consists of several blue line segments that intersect to form a network of polygons. There are eight red square points scattered across the diagram, each representing a seed point for one of the Voronoi regions. The lines extend from the vertices towards the edges of the frame.

Voronoi Diagrams

Advanced Topics and Metrics

Chapters 17 and 18

Definitions

- **M** : a set of n points in \mathbb{E}^d , M_1, \dots, M_n , called sites.

$$V(M_i) = \{X \in \mathbb{E}^d : \delta(X, M_i) \leq \delta(X, M_j), j \neq i\}$$


- $V(M_i)$: The intersection of a finite number of closed half-spaces, perpendicular bisectors of point pairs.
- **$\text{Vor}(M)$** : A complex of $V(M_i)$ cells.
- L_2 -general position: no sphere contains $d+2$ sites on its boundary

Diagram as regions of Sphere Centers

- $V(M_i)$ as the set of the center of spheres
 - Boundary of contains M_i , interior contains no M_j .
 - When $d=2$
 - Voronoi Edges: centers of spheres containing exactly 2 points on their edge
 - Voronoi Verteces: center of spheres containing 3 or more points on the edge (in L_2 -general, exactly 3).

Voronoi Diagrams and Polytopes


- By treating points as spheres, we can describe the diagram as a polytope in $d+1$ dimensions



Power of a point w.r.t. a sphere


- \mathbb{E}^d : D dimension Euclidean space, origin O .
- X : a point in \mathbb{E}^d .
- Σ : a sphere in \mathbb{E}^d with center C , radius r .
 $\Sigma(X)=0$, where: $\Sigma(X)=|XC|^2-r^2$ (17.1)
- Interior of the sphere: $\text{int}(\Sigma)=\{X: \Sigma(X)<0\}$
- Exterior of the sphere: $\text{ext}(\Sigma)=\{X: \Sigma(X)>0\}$

- σ : Power w.r.t. origin: $\Sigma(O) = \sigma = |C|^2 - r^2$ (17.2)
- If D is any line containing X , and intersecting Σ at A and B $\Sigma(X) = |XA| \cdot |XB|$ (17.3)
- When D is tangent at T , $\Sigma(X) = XT^2$ (17.4)



Representation of Spheres

- Let ϕ be a mapping that takes a sphere Σ in \mathbb{E}^d with center C , power w.r.t. O σ to the point $\phi(\Sigma)=(C, \sigma)$ in \mathbb{E}^{d+1} .
- Allows us to treat spheres in \mathbb{E}^d as points in \mathbb{E}^{d+1}



Embedding Spheres as Points

- Embed \mathbb{E}^d as the hyperplane in \mathbb{E}^{d+1} , with $x_{d+1}=0$.
 - x_{d+1} is the vertical direction, positive is up
 - Each point in \mathbb{E}^d is a sphere of radius $|C|$.
- X denotes a point in \mathbb{E}^d with vector (x_1, \dots, x_d) ,
- \underline{X} denotes a point in \mathbb{E}^{d+1} , with vector (x_1, \dots, x_{d+1}) .
- Projection means a vertical projection.

Homogeneous Coordinates and Matrix Notation

- $\mathbf{X}=(x_1,\dots,x_d,t)$, $\underline{\mathbf{X}}=(x_1,\dots,x_{d+1},t)$
- We can then rewrite the equation of the sphere, Σ , as $\mathbf{X}\underline{\Sigma}\underline{\mathbf{X}}^t = 0$, with

$$\underline{\Sigma} = \begin{pmatrix} I_d & -C^t \\ -C & \sigma \end{pmatrix}$$

Paraboloid P


- ϕ maps points in \mathbb{E}^d , as spheres of radius 0, to a paraboloid of revolution, P , with vertical axis and equation:

$$x_{d+1} = \sum_{i=1}^d x_i^2 = X \cdot X$$

- In homogenous coordinates:

$$\underline{X} \Delta_P \underline{X}^t = 0$$

$$\Delta_P = \begin{pmatrix} I_d & 0 & 0 \\ 0 & 0 & -1/2 \\ 0 & -1/2 & 0 \end{pmatrix}$$



Paraboloid P

- Real spheres: non-negative radius; on or below this paraboloid
- Imaginary spheres: negative radius; above this paraboloid.

Polarity

- Any hyperplane H in a projective space has a homogeneous equation of the form

$$H = \{X : \sum_{i=1}^{d+1} h_i x_i = 0\}$$

- Let S be the matrix $S = \begin{pmatrix} I_d & 0 \\ 0 & -1 \end{pmatrix}$
- We let H^* be the point $(h_1, \dots, h_d, -h_{d+1})$.

$$H = \{X : H^* * S X^t = 0\}$$

Polarity

- Chapter 7.3
- The *polar set*, A^* : Given a point A ,
 $A^* = \{X : \langle X, A \rangle = 1\}$
 - A^* the polar hyperplane of A
 - A is the pole of A^*
- Polarity is the duality that connects points and planes.

Polarity w.r.t. Q

- Consider the quadric Q in \mathbb{E}^{d+1} , defined by the homogeneous equation: $Q(\underline{X}) = \underline{X}\Delta_Q\underline{X}^t = 0$
- Points X and Y are conjugates w.r.t. Q if
$$Q(\underline{X}, \underline{Y}) = \underline{X}\Delta_Q\underline{Y}^t = 0$$
- $A^* = \{ \underline{X} : \underline{A}\Delta_Q\underline{X}^t = 0 \},$
- Polarity w.r.t. Q maps hyperplane \underline{H} to the pole \underline{H}^*

Polarity w.r.t. P

- $A^* = \{ \underline{X} : \underline{A} \Delta_Q \underline{X}^t = 0 \}$
- If $Q=P$, and $\phi(\Sigma)=(C,\sigma)$, $\phi(\Sigma)^*$ can be rewritten as $x_{d+1} = 2C \cdot X - \sigma$
- Since polarity preserves incidences: for point \underline{X} and hyperplane \underline{H} ,

$$\underline{X} \in \underline{H} \text{ iff } \underline{H}^* \in \underline{X}^*$$

- This also preserves the side of the parabola

Orthogonal Spheres

- Σ_1, Σ_2 (centers and radii C_1, C_2, r_1, r_2) are *orthogonal* if

$$\Sigma_1(C_2) = r_2^2$$

$$\Sigma_2(C_1) = r_1^2$$


- We can rewrite these equivalent equations as

$$C_1 \cdot C_2 - \frac{1}{2}(\sigma_1 + \sigma_2) = 0$$

- This shows that 2 spheres are orthogonal if their mappings are conjugate w.r.t. P .

Lemma 17.2.1

- The set of spheres in \mathbb{E}^d that are orthogonal to a given sphere is mapped by ϕ to the polar hyperplane $\phi(\Sigma)^*$




Paraboloid and Hyperplane

- The set of spheres that passing through X = the set of spheres that intersect sphere with $r=0$, centered at X .
- Image of these spheres is $\phi(X)^*$
- $\phi(X)^*$ is tangent to P , since X is only sphere with $r=0$ orthogonal to X

Lemma 17.2.2

- Let Σ be a sphere in E^d .
- The intersection of $\phi(\Sigma)^*$ with P is the image under ϕ of the set of spheres with radius 0 that are orthogonal to Σ , namely Σ itself
- $\phi(\Sigma)^* \cap P$ in E^{d+1} projects onto Σ in E^d .



Lemma 17.2.2


- $P \cap \underline{H}$ projects onto \mathbb{E}^d as a sphere $\phi^{-1}(\underline{H}^*)$, centered at the vertical projection of \underline{H}^*
- The points of sphere Σ lifted onto P in \mathbb{E}^{d+1} belong to a unique hyperplane that intersects P exactly at these points. This hyperplane is $\phi(\Sigma)^*$

Lemma 17.2.3

- From 17.2.2, the power of X w.r.t. Σ equals square of the radius of the sphere Σ_X
 - Orthogonal to Σ , centered at X
- This can be computed in \mathbb{E}^{d+1}
- Σ_X is mapped to a point \underline{l} in \mathbb{E}^{d+1} , the intersection of the vertical line passing through X with $\phi(\Sigma)^*$
- The x_{d+1} coordinates of $\phi(X)$ and \underline{l} are:
 X^2 and $\Sigma_X(O) = X^2 - \Sigma_X(X)$, because $r_{\Sigma_X} = \Sigma(X)$

Lemma 17.2.3

- The power of X with respect to a sphere Σ equals the signed vertical distance from the point $\phi(X)$ to the hyperplane $\phi(\Sigma)^*$



Lemma 17.2.4

- Let X and Σ be respectively a point and a sphere in \mathbb{E}^d . If \underline{H} is a hyperplane in \mathbb{E}^{d+1} , we denote by \underline{H}^- the half-space lying below \underline{H} . Then:

$$X \in \Sigma \Leftrightarrow \phi(X) \in \phi(\Sigma)^* \Leftrightarrow \phi(\Sigma) \in \phi(X)^*$$

$$X \in \text{int}(\Sigma) \Leftrightarrow \phi(X) \in \phi(\Sigma)^{*^-} \Leftrightarrow \phi(\Sigma) \in \phi(X)^{*^-}$$

$$X \in \text{ext}(\Sigma) \Leftrightarrow \phi(X) \in \phi(\Sigma)^{*^+} \Leftrightarrow \phi(\Sigma) \in \phi(X)^{*^+}$$

$$X \in \Sigma \Leftrightarrow \phi(X) \in \phi(\Sigma)^* \Leftrightarrow \phi(\Sigma) \in \phi(X)^*$$

$$X \in \text{int}(\Sigma) \Leftrightarrow \phi(X) \in \phi(\Sigma)^{-*} \Leftrightarrow \phi(\Sigma) \in \phi(X)^{-*}$$

$$X \in \text{ext}(\Sigma) \Leftrightarrow \phi(X) \in \phi(\Sigma)^{+*} \Leftrightarrow \phi(\Sigma) \in \phi(X)^{+*}$$

- Any point lying below $\phi(X)^*$ is the image of a sphere whose interior contains X . Any point above is the image of a sphere whose exterior contains X .

Radical Hyperplanes

- Given Σ_1, Σ_2 , the set of points in \mathbb{E}^d with the same power with respect to Σ_1 and Σ_2 is the *radical hyperplane* H_{12} :

$$H_{12} : \Sigma_1(X) - \Sigma_2(X) = 0$$


- The points that are the centers of spheres orthogonal to both spheres
- These points are mapped to $\phi(\Sigma_1)^* \cap \phi(\Sigma_2)^*$

Application to Voronoi Diagrams

- By 17.2.1, $\phi(M_i)^*$ denotes the polar hyperplane tangent to P at $\phi(M_i)$
- The set of spheres containing NO site M_i is mapped to the intersection of the half-spaces lying above the hyperplanes


Voronoi Polytope

- The aforementioned intersection is an unbounded polytope that contains P .
- This is the *voronoi polytope*, $V(M)$



Theorem 17.2.5

- The Voronoi Diagram of \mathbf{M} , $\mathbf{Vor}(\mathbf{M})$, is a cell complex of dimension d in \mathbb{E}^d whose faces are obtained by projecting onto \mathbb{E}^d the proper faces of the Voronoi Polytope $\mathbf{V}(\mathbf{M})$



Proof of 17.2.5

- The boundary of $V(\mathbf{M})$ is a pure cell complex of dimension d , hence so is $\mathbf{Vor}(\mathbf{M})$.
- Let A be a point on a facet of $V(\mathbf{M})$, that is contained in $\phi(M_i)^*$.
- A is the image of a sphere centered on the projection of A that passes through M_i , and contains no other site.
- Therefore, A must belong to $V(M_i)$


Implications

- Combinatorial properties of VDs follow directly from those of polytopes
 - If points are in general position, $V(\mathbf{M})$ is a simple $(d+1)$ -polytope
 - Each vertex is incident to $d+1$ hyperplanes
 - This is the L_2 -*general* assumption
- Computing the VD of n points \approx computing the intersection of n half-spaces in 1 higher dimension

Corollary 17.2.6

- The complexity (# of faces) of Voronoi Diagrams of n points is $\Theta(n^{\text{ceil}(d/2)})$
- This can be computed in $O(n \log n + n^{\text{ceil}(d/2)})$ time, which is worst case optimal

Delaunay Complexes



Delaunay Complexes


- Define \mathbf{M} as before, and the mapping onto the paraboloid P as well.
- Denote by $\mathbf{D}(\mathbf{M})$ the convex hull of the mappings of the points M_i as well as some 'high' point O' .
- K -face: a k -dimensional face:
 - 0-face is vertex, 1-face is a line, $d-1$ is a facet

$D(M)$

- $D(M)$ forms the convex hull of $\phi(M)$ and some point O' , a high point on x_{d+1} axis, such that the convex hull is stable as O' rises to infinity
- The faces of this convex hull that do not contain O' form the lower envelope of $\text{conv}(\phi(M))$
- $\text{Del}(M)$ is the projection of the convex hull.

$Del(M)$

- A vertical projection of the polytope $D(M) \subseteq \mathbb{E}^d$.
- The k -faces of $Del(M)$ are in 1-to-1 correspondence with the k -faces of $D(M)$ that do not contain O'



Vor(M) and ***Del(M)***

- Exists a bijection between the faces of ***V(M)*** and the faces of ***D(M)*** that do not contain O'
 - Maps the facet of ***V(M)*** containing $\phi(M_i)^*$ to the point $\phi(M_i)$.
- The bijection exists between the k -faces of ***Vor(M)*** and the $(d-k)$ -faces of ***Del(M)*** that reverses inclusion relationships
- ***Del(M)*** is the dual of ***Vor(M)***.

Theorem 17.3.1

- The Delaunay complex of \mathbf{M} is the dual to the Voronoi Diagram.
- Its faces are obtained by projecting the faces of the lower envelope of the convex hull of $\phi(\mathbf{M})$, obtained by lifting the \mathbf{M} onto P .
- Computing the Delaunay complex is equivalent to computing the Convex hull

Corollary 17.3.2

- The Delaunay complex of n sites in \mathbb{E}^d can be computed in time $O(n \log n + n^{\text{ceil}(d/2)})$ time, which is worst case optimal

Delaunay Triangulations

- In L_2 -general, $V(\mathbf{M})$ is a simple polytope, $D(\mathbf{M})$ is a simplicial polytope, and $Del(\mathbf{M})$ is a simplicial complex.
- If not L_2 -general, some of the faces are not simple (triangles).
 - There could be many valid triangulations. All of them are considered Delaunay triangulations.

Characteristic Properties

- Theorem 17.3.3
 - Any d -face in the complex can be circumscribed by a sphere that passes through all its vertices M_{i0}, \dots, M_{il} , and whose interior contains no site of M

Proof of 17.3.3

- \mathbf{M}_k : a subset of \mathbf{M} containing k sites
- Pick a d -face, T , of $\mathbf{Del}(\mathbf{M})$, with vertices \mathbf{M}_k . T is the convex hull of \mathbf{M}_k . The CH of $\phi(\mathbf{M}_k)$ form a d -face F of the total CH, by thm 17.3.1.
- \underline{H}_F : the hyperplane that supports F
- $\underline{H}_F \cap P$ projects onto \mathbb{E}^d as a sphere Σ circumscribed to $\text{conv}(M_{i0}, \dots, M_{il})$, centered on the projection of the pole, \underline{H}_F^* .

Proof, Continued

- \underline{H}_F^* is the intersection of the polar hyperplanes $\phi(\mathbf{M}_k)$, and it is projected onto \mathbb{E}^d at C
- C is the vertex of the VD that is incident to the cells that correspond to the sites \mathbf{M}_k , and none of the interiors can contain any other site.

Theorem 17.3.4

- Let \mathbf{M}_k be a subset of k sites in \mathbf{M} . The CH of \mathbf{M}_k is a face of the Delaunay complex iff there exists a $(d-1)$ sphere passing through the vertexes \mathbf{M}_k , and such that no point in \mathbf{M} is on the interior of this sphere.

Proof of 17.3.4

- Necessary
 - Result of 17.3.3, and sphere circumscribed to a face also circumscribed to a subface.
- Assume exists a $(d-1)$ -sphere Σ that passes through the points of M_k , and has no interior sites. Let \underline{H} be the hyperplane, $\phi(\Sigma)^*$ containing the projections of the points

Proof of 17.3.4

- The halfspace lying below \underline{H} does not contain any points in $\phi(\mathbf{M})$, according to 17.2.4
- Thus, \underline{H} is a hyperplane supporting $\mathbf{D}(\mathbf{M})$ along the convex hull of the k sites, and so $\text{conv}(\phi(M_{i_0}) \dots \phi(M_{i_k})) = \underline{H} \cap \mathbf{D}(\mathbf{M})$ is a face of $\mathbf{D}(\mathbf{M})$
- Therefore, from 17.3.1, \mathbf{M}_k makes a face of $\mathbf{Del}(\mathbf{M})$

Corollary 17.3.5

- Any Delaunay Triangulation of a set of **M** sites is such that the sphere circumscribed to any d -simplex in the triangulation contains no point of M in its interior.
- Conversely, any triangulation satisfying this property is a Delaunay Triangulation

Characteristic of DT

- Consider any DT $\mathbf{T}(\mathbf{M})$
- Let $S_1 = M_1 \dots M_d M_{d+1}$ and $S_2 = M_1 \dots M_d M_{d+2}$ be a pair of adjacent d -simplices in $\mathbf{T}(\mathbf{M})$ (circumscribed to $\Sigma_1 \Sigma_2$). that share a common face $F = M_1 \dots M_d$
- (S_1, S_2) is regular if M_{d+1} is not in $int(\Sigma_2)$

Characteristics

- If Σ_1 differs from Σ_2 , regularity is equivalent to M_{d+2} not being in $int(\Sigma_1)$.
- M_{d+1} does not belong to $int(\Sigma_2)$ iff $\Sigma_2(M_{d+1}) > 0$.
But the hyperplane \underline{H}_F that supports F is the radical hyperplane of Σ_1 and Σ_2

Theorem 17.3.6

- Since $\Sigma_1(M_{d+1})=0$, the half-space bounded by \underline{H}_F that contains M_{d+1} (or resp. M_{d+2}) consists of the points whose power w.r.t. Σ_1 is smaller (resp. greater) than their power w.r.t. Σ_2 .

$$\Sigma_1(M_{d+2}) > \Sigma_2(M_{d+2})=0$$

- This proves M_{d+2} does not belong to the interior of Σ_1

Theorem 17.3.6

- Consider a triangulation $\mathbf{T}(\mathbf{M})$
- Then, $\mathbf{T}(\mathbf{M})$ is a DT iff all pairs of adjacent d -simplices in $\mathbf{T}(\mathbf{M})$ are regular

Proof of 17.3.6

- Necessary as a consequence of 17.3.3
- $\phi(S)$ = the k -simplex whose vertices are the images of the vertices of a k -simplex S .
- C = the union of the $\phi(S)$'s for all the faces S of the DT $T(\mathbf{M})$.
- The sufficiency proof is to show that C is the graph of a real valued convex function over $\text{conv}(\mathbf{M})$

Proof, Cont.

- We consider S_1 and S_2 that share the common face F , with circumscribing spheres Σ_1 and Σ_2
- By 17.2.4, the regularity condition is equivalent to $\phi(M_{d+1})$ being in $\phi(\Sigma_2)^{*+}$, and vice versa.
- If (S_1, S_2) regular, $\phi(F)$ is locally convex (there is a hyperplane containing $\phi(F)$ such that $\phi(S_1)$ and $\phi(S_2)$ belong to the half-space above this hyperplane)

Proof Cont

- This is true for any $(d-1)$ -face of C incident to 2 d -faces, and so C is locally convex at any point
- C is defined over a convex subset, the $\text{conv}(\mathbf{M})$.
- Therefore, C is convex and is the lower envelope of the polytope $\mathbf{D}(\mathbf{M})$ which proves that $\mathbf{T}(\mathbf{M})$ is a DT of \mathbf{M} .

Optimality of a DT

- Because we can triangulate in many ways, we have many ways to define optimality
 - Compactness
 - Equiangularity

Compactness

- The smallest enclosing sphere of each simplex S
- $T(\mathbf{M})$ corresponds to a function $\Sigma_T(\mathbf{M})$, defined over $\text{conv}(\mathbf{M})$ as the power of a point X w.r.t. the sphere Σ circumscribing any d -simplex containing X

Lemma 17.3.7

- Let ***Det***(***M***) be a Delaunay triangulation of ***M***, and ***T***(***M***) be another triangulation. Then
 - For all X in $\text{conv}(\mathbf{M})$, $\Sigma_{\text{Det}}(X) \geq \Sigma_T(X)$

Proof of 17.3.7

- Consider a d -simplex T in $\mathcal{T}(\mathbf{M})$ containing X , Σ the circumscribed sphere, and $\phi(T)$ the d -simplex projection
- By 17.2.3, $\Sigma_T(X)$ is the (negative) vertical distance from $\phi(\Sigma)^*$ to $\phi(X)$.
- $\phi(\Sigma)^*$ is the affine hull of $\phi(T)$
- For a given X , this signed vertical distance is maximized when $\phi(T)$ is a face of the convex hull of $\phi(\mathbf{M})$; when T is a simplex of a DT

Lemma 17.3.8

- If T is a d -simplex circumscribed to Σ_T , then

$$\min_{X \in T} \Sigma_T(X) = \Sigma_T(C'_T) = -r'^2_T$$

- where C'_T and r'_T are the center and radius of the smallest enclosing sphere

Proof

- Let Σ_T be the sphere circumscribed to T , centered at C_T with radius r_T .
 - $\Sigma_T(X) = XC_T^2 - r_T^2$ is minimized when $X = C_T$, and so $\Sigma_T(X)$ is greater than $-r_T^2$.
- If C_T contained in T , Σ_T is the smallest enclosing sphere, and so $r'_T = r_T$

Proof, Cont.

- If C_T not contained in T , the smallest containing sphere is centered on a k -face ($k < d$) F , the face such that the orthogonal project of C_T onto the plane that supports F falls inside F .
- r'_T of this sphere is that of the $(k-1)$ -sphere circumscribed to F
- C'_T minimizes XC_T when X is in T
- $C_T C'_T + r'^2_T = r^2_T$

Most Compact Triangulation

- The *maximum min-containment radius* of $\mathcal{T}(\mathcal{M})$: $C(\mathcal{T}(\mathcal{M})) = \max_{T \in \mathcal{T}(\mathcal{M})} r'_T$
- The most compact triangulation minimizes $C(\mathcal{T}(\mathcal{M}))$.

Theorem 17.3.9

- Delaunay Triangulations are the most compact among all triangulations
 - Other triangulations might also be the most compact, even if not Delaunay

Proof of 17.3.9


- Define $T(\mathbf{M})$, $Det(\mathbf{M})$
- X_T is the point minimizing $\Sigma_T(X)$, and similarly X_{Det} .
- X_T is the center of the smallest sphere circumscribed to the simplex containing X_T (by 17.3.8)
- We denote the radii as r'_T and r'_{Det} .
- $C(T(\mathbf{M})) = r'_T$, $C(Det(\mathbf{M})) = r'_{Det}$
- By 17.3.7 and 8:
- $\Sigma_T(X_T) = -r'^2_T \leq \Sigma_T(X_{det}) \leq \Sigma_{Det}(X_{Det}) = -r'^2_{Det}$

Equiangularity (d=2)

- Given a triangulation $T(\mathbf{M})$, the angle vector is $Q(T(\mathbf{M})) = (a_1, \dots, a_{3t})$, where each a is an angle of the t triangles, sorted by increasing value.
- We know that the sum of the angles $= t\pi$
- A triangulation that maximizes the angle vector for the lexicographic ordering also maximizes the smallest ordering
- This is a *globally equiangular* triangulation

Theorem 17.3.10


- A globally equiangular triangulation of a set of M sites in the plane is always a DT



Proof of 17.3.10



- Consider 2 triangles $T_1=ABC$, and $T_2=BCD$, where $ABCD$ is a strictly convex quadrilateral (all vertices of the convex hull)
- We can flip diagonals to increase equiangularity
 - This is a regularization rule, since it changes a pair of adjacent triangles into a regular pair of triangles

- Let Σ_1 and Σ_2 be the circumscribing circles.
 - AD is flipped only if Σ_1 contains D
- Let $abcd$ be the angles at the vertices ABCD, and b and c are split into b_1, b_2, c_1 and c_2
 - We also split a and d into a'_1, a'_2, b'_1, b'_2 for the potential split



- If a is the smallest angle, then we don't flip, but $d = \pi - b_2 - c_2 < \pi - a$, so that $a + d < \pi$, so A is not inside Σ_2
- If b_1 is the smallest, then we flip only if d'_2 is greater, which only happens when D in Σ_1
- Parallels for all other smallest angles
- After we flip, $Q(T_1(M)) > Q(T(M))$
- Progressive flippings increases the angle vector. Since there are a finite number of triangulations, we must reach a maximum, with only regular pairs of adjacent triangles, and is a DT by 17.3.6

Higher Order Voronoi Diagrams



Higher Order Voronoi Diagrams

- Level- k
 - A point is at Level k of an arrangement \mathbf{A} if it belongs to exactly k open half-spaces, $\phi(M_i)^*$, such that each half-space does not the reference point (the origin).

Voronoi Diagram of Order k

- Let \mathbf{M}_k be a subset of \mathbf{M} of size k .
- $V_k(\mathbf{M}_k)$: the region or points that are closer to \mathbf{M}_k than any other sites

$$V_k(M_k) = \left\{ X : \forall M_i \in M_k, \forall M_j \in M \setminus M_k, \|XM_i\| \leq \|XM_j\| \right\}$$

- The union of V_k forms $\mathbf{Vor}_k(\mathbf{M})$
- $\mathbf{Vor}_1(\mathbf{M}) = \mathbf{Vor}(\mathbf{M})$.

Theorem 17.4.1

- **$\text{Vor}_k(\mathbf{M})$** is a cell complex of dimension d in \mathbb{E}^d . The cells of this complex correspond to the cells at level k in the arrangement **\mathbf{A}** of the hyperplanes induced by the projections of the points. A cell in the diagram is obtained by projecting the corresponding cell in the arrangement. The l -faces of **$\text{Vor}_k(\mathbf{M})$** are obtained by projecting the l -faces common to cells at level k in arrangement **\mathbf{A}** .

Proof of 17.4.1



- The proof relies on lemma 17.2.4
- A sphere whose interior contains k points is mapped by ϕ to a point at level k in the arrangement **A** of the hyperplanes.
- This is easily verifiable.

$Vor_k(M)$


- Once constructed, easy to find k -nearest neighbors, by finding the region of the arrangement containing the point.
- Theorem 17.4.2
 - The overall complexity of the first k voronoi diagrams of a set of n points is $O(n^{\text{ceil}((d+1)/2)} k^{\text{ceil}((d+1)/2)})$. The k diagrams can be computed in time $O(n^{\text{ceil}((d+1)/2)} k^{\text{ceil}((d+1)/2)})$ for $d \geq 3$, or $O(nk^2 \log(n/k))$ for $d=2$

Examples

- $\text{Vor}_{n-1}(\mathbf{M})$ is the furthest-point Voronoi Diagram.
- $\text{Vor}_2, \text{Vor}_3$



Projection of $V_{\text{ork}}(M)$



Non-Euclidean Metrics

Power Diagrams


- Let $S = \{\Sigma_1, \dots, \Sigma_n\}$ be a set of spheres in E^d .
- $P(\Sigma_i) \rightarrow$ all points whose power w.r.t. Σ_i is smaller than power w.r.t. any other sphere.

$$P(\Sigma_i) = \{X \in E^d : \forall j \neq i, \Sigma_i(X) \leq \Sigma_j(X)\}$$

- The region $P(\Sigma_i)$ is a convex polytope, the intersection of a finite set of half-spaces bounded by radical hyperplanes.
- The complex is called the Power Diagram of S , $Pow(S)$.

Power Diagrams

- As before, we map the spheres to a point in \mathbb{E}^{d+1} . However, these points are not on the paraboloid. The faces of the Power diagram are obtained by projecting the intersections of the polar hyperplanes.



Order- k Power Diagrams

- A similar parallel to the previous.

Affine Diagrams

- Defined for sites and for a distance such that the set of points equidistant from 2 objects is a hyperplane.
- We can extend the idea of VD's to more general sites, non-Euclidean distances.
- VD's and Power Diagrams are *Affine Diagrams*, and any affine diagram is a power diagram.
- We can also derive many nonaffine diagrams from affine diagrams

Affine Diagrams

- Cells of AD's are convex polytopes.
- To any AD of n objects corresponds a set of $\binom{n}{2}$ perpendicular bisectors, H_{ij} , $1 \leq i < j \leq n$
- These hyperplanes must satisfy
$$H_{ij} \cap H_{jk} = H_{ij} \cap H_{ik} = H_{ik} \cap H_{jk} \stackrel{\text{def}}{=} I_{ijk}$$
- For any $1 \leq i < j < k \leq n$
- The diagram is simple if the I_{ijk} are disjoint and nonempty

Theorem 18.2.1

- Any simple AD in \mathbb{E}^d is the power diagram of a set of spheres in \mathbb{E}^d .
- Proof:
 - We embed \mathbb{E}^d in \mathbb{E}^{d+1} as the hyperplane of $x_{d+1}=0$.
 - We construct a set of n hyperplanes in \mathbb{E}^{d+1} such that the projection of the intersections is exactly the hyperplanes H_{ij}

Proof of 18.2.1

- We construct n hyperplanes, P_1, \dots, P_n , in \mathbb{E}^{d+1}
- We ensure that the vertical projections of $P_i \cap P_j$ for any $i < j$ is exactly H_{ij}
- Each P_i is the polar hyperplane of a sphere $\Sigma_i = \phi^{-1}(P_i^*)$, which is the projection of $P_i \cap P$.
- So H_{ij} is the radical hyperplane of Σ_i and Σ_j . So the affine diagram is the power diagram

Constructing P_i

- Denote h_{ij} the vertical projection of H_{ij} onto P_i .
- Take P_1 and P_2 as non-vertical hyperplanes that intersect P_1 along h_{12} . For any $k > 2$, we must take P_k that make the appropriate intersections with P_1 and P_2 . If l_{12k} exists, then P_k must exist.

Proof Cont.

- We show that the projections of the intersections are H_{ij} .
- We have constructed H_{12} , H_{1k} , and H_{2k} for one k .
- If the other cells I_{ijk} exist, then their intersections must exist

Theorem 18.2.2

- The affine diagram with Hyperplanes H_{ij} and equations of hyperplanes
$$-2(C_i - C_j) \cdot X + \sigma_i - \sigma_j = 0$$
is the power diagram of the spheres Σ_i .
- Proof: We can rewrite the equation of the hyperplane as $\Sigma_i(X) - \Sigma_j(X) = 0$

Diagrams for General Quadratic Distance

- Consider 2 points X and A in \mathbb{E}^d .
- The General Quadratic Distance is defined as:
$$\delta_Q(X, A) = (X - A)\Delta(X - A)^t + p(A)$$
- Where Δ is a real symmetric $d \times d$ matrix and where $p(A)$ is a real number

Diagrams

- VD is when Δ is the identity matrix and $p(X)=0$
- Furthest point diagrams (VD of order $n-1$) are when Δ is the negative identity matrix, and $p(X)=0$.
- Power Diagrams are when Δ is the identity matrix and $p(X)$ does not equal 0.

General Quadratic Distance

- If A, B are points, then the formula for the hyperplane can be rewritten as

$$H_{AB} : 2(B - A)\Delta X^t + A\Delta A^t - B\Delta B^t + p(A) - p(B) = 0$$

- Therefore, the VD for a general quadratic distance is an affine diagram

Weighted Diagrams


- Non-affine.
- Defined over finite sets of points and a weighted Euclidean distance

Additively Weighted Diagrams: $Vor_+(M)$

- Additive weighted distance formula


$$\delta_+(X, M_i) = \|XM_i\| - r_i$$

- We assume that all r_i values, the weights, are non-negative.



$Vor_+(M)$

- Consider Σ_i in \mathbb{E}^d , centered at M_i with radius r_i , and let ψ be the bijection that maps Σ_i to the point $\psi(\Sigma_i)=(M_i, r_i)$ in \mathbb{E}^{d+1}
- The spheres of radius 0 correspond to the hyperplane $x_{d+1}=0$ in \mathbb{E}^{d+1} .



Projection ψ

- Points at additive distance r from M_i are centers of spheres tangent to Σ_i with radius r , inside or outside Σ_i .
- Ψ generates a cone of revolution
$$C(\Sigma) : x_{d+1} = |XC| - r$$
 - Apex of C is $(C, -r)$,
 - symmetrical with respect to $x_{d+1} = 0$,
 - has an aperture angle $\pi/4$

$$\psi$$


- The projection I_X of a point X on the cone $C(\Sigma)$ is the projection under ψ of the sphere centered at X and tangent to Σ
- Signed vertical distance from X to I_X equals the additive distance from X to C weighted by r
- Each sphere, Σ_i , corresponds to a cone $C(\Sigma_i)$, also denoted C_i .
- The projection of the lower envelope of the cones is exactly **$\text{Vor}_+(M)$**

Equidistance w.r.t. additive distance

- The set of points in \mathbb{E}^d equidistant (w.r.t. additive distance) from 2 points == the projection of intersection of cones
- $C_1 : (x_{d+1}+r_1)^2 = XM_1^2, x_{d+1}+r_1 > 0$
- $C_2 : (x_{d+1}+r_2)^2 = XM_2^2, x_{d+1}+r_2 > 0$
- $H_{12} : -2(M_1-M_2) \cdot X - 2(r_1-r_2)x_{d+1} + M_1^2 - r_1^2 - M_2^2 + r_2^2$

\mathbf{Vor}_+ and Power Diagram

- There exists a correspondence between \mathbf{Vor}_+ and a power diagram
- Take spheres Σ_i' in \mathbb{E}^{d+1} , centered at $\psi(\Sigma_i)$ with radius $r_i\sqrt{2}$



Vor_+ and PD

- The cell of **$Vor_+(M)$** that corresponds to M_i , $V_+(M_i)$, is the projection of intersection of the cone C_i with the cell of the PD corresponding to the sphere Σ_i'
- X is in $V_+(M_i)$ iff the projection X_i of X onto C_i has a smaller x_{d+1} coordinate than of $X_j, j \neq i$
 (X, x_{d+1}) of X_i must obey:
 $(x_{d+1} + r_i)^2 = XM_i^2, (x_{d+1} + r_j)^2 \leq XM_j^2, j \neq i$

Computation of \mathbf{Vor}_+

- $\Sigma_i'(X_i) \leq \Sigma_j'(X_i)$ for any $j \neq i$
- \mathbf{Vor}_+ can be computed by
 - Compute Σ_i'
 - Compute the power diagram of the Σ_i' s
 - For all i , project onto \mathbb{E}^d the intersection of C_i with the cell of the PD that corresponds to Σ_i' .

Computation of \mathbf{Vor}_+

- The PD of Σ_i' can be computed in time $O(n^{\lfloor d/2 \rfloor + 1})$
- The intersections can be computed in $O(n^{\lfloor d/2 \rfloor + 1})$
- Theorem 18.3.1
 - \mathbf{Vor}_+ has complexity $O(n^{\lfloor d/2 \rfloor + 1})$
 - Can be computed in time $O(n^{\lfloor d/2 \rfloor + 1})$


Optimality of Computation

- This result is optimal in odd dimensions, but is not optimal in $d=2$, and possibly not for any even dimension (can be computed in optimal time $O(n \log n)$)

Multiplicatively Weighted Diagrams

$\mathbf{Vor}_*(M)$

- Similar to \mathbf{Vor}_+ , but with multiplicative distance $\delta_*(X, M_i) = p(M_i) \|XM_i\|$
- Where $p(M_i)$ is a positive real number (from now on, p_i)



Vor*

- Set of points at equal multiplicative distance from 2 sites M_i, M_j is a sphere of the equation
$$p_i(X-M_i)^2 = p_j(X-M_j)^2$$

- In normalized form

$$X^2 - 2 \frac{p_i M_i - p_j M_j}{p_i - p_j} \cdot X + \frac{p_i M_i^2 - p_j M_j^2}{p_i - p_j} = 0$$

- And in \mathbb{E}^{d+1} , the sphere can be represented as

$$\phi(\Sigma_{ij}) = \left(\frac{p_i M_i - p_j M_j}{p_i - p_j}, \frac{p_i M_i^2 - p_j M_j^2}{p_i - p_j} \right)$$

Point $\phi(\Sigma_{ij})$

- We have mapped the sphere intersection to a point in \mathbb{E}^{d+1} .
- The polar hyperplane H_{ij} , w.r.t. the paraboloid P has equation

$$\begin{aligned} H_{ij}(X, x_{d+1}) = \\ (p_i - p_j)x_{d+1} - 2p_i M_i \cdot X + 2p_j M_j \cdot X + p_i M_i^2 - p_j M_j^2 \\ = 0 \end{aligned}$$

$$\phi(\Sigma_{ij})$$

- H_{ij} 's are the radical hyperplanes of spheres Σ_i in \mathbb{E}^{d+1}
- Σ_i is centered at $(p_i M_i, -\pi/2)$, with $\sigma = p_i M_i^2$.
- We now have a correspondence between ***Vor**** and the PD of Σ_i 's

***Vor*_{*}** and PD

- X is a point in \mathbb{E}^d , projected onto P

$$X \in V_*(M_i) \Leftrightarrow p_i(X - M_i)^2 \leq p_j(X - M_j)^2$$

$$\Leftrightarrow H_{ij}(X, X^2) \leq 0$$

$$\Leftrightarrow \Sigma_i(\phi(X)) \leq \Sigma_j(\phi(X))$$

$$\Leftrightarrow \phi(X) \in P(\Sigma_i)$$

Computing ***Vor****

- Compute Σ_i
- Compute PD of the Σ_i 's
- For each I , project the intersection of the Σ_i in the PD with the paraboloid P

Complexity of computation

- Theorem 18.3.2
 - ***Vor***_{*} has complexity $O(n^{\lfloor d/2 \rfloor + 1})$ and can be computed in time $O(n^{\lfloor d/2 \rfloor + 1})$
- This is optimal

L_1 metric


- Reminder, L_1 distance between X and point M

$$\delta_1(X, M) = \sum_{i=1}^d |x_i - m_i|$$

- Points at distance r from M are a polytope w/ vertices at coordinates $x_i = m_i \pm r$, and $x_i = m_i$ if $i \neq j$

- In 2d- a tilted square

- In 3d- a regular octahedron




- These are dual to the cube, a co-cube

Vor_{L1}

- Voronoi Diagram for L_1 distance
- We create a similar mapping to **Vor_+** , but mapped to a pyramid instead of a cone, with

$$x_{d+1} = \delta_1(X, M_i)$$




Vor_{L1}

- Lower envelope of the Pi pyramids
 - The graph of the function $\min_{1 \leq i \leq n} \delta_1(X, M_i)$
- Each portion from a distinct pyramid projects onto \mathbb{E}^d as a facet, a cell of **$Vor_{L1}(M)$**
- The complexity is $O(n^d a(n))$.

Region Bisectors in Vor_{L1}


- When $d=2$, bisectors are generally 3 piece polylines. However, can be 1 linear segment connecting 2 regions of dimension 2



L_∞ Metric

- Using the distance metric

$$\delta_\infty(X, M) = \max |x_i - m_i|$$




Voronoi Diagrams in Hyperbolic Spaces


Pencils of spheres

- A set of spheres, S , that are affine combinations of 2 given spheres, Σ_1 and Σ_2
 - $F = \{\Sigma : \text{exists real } \lambda, \text{ s.t. for all } X \text{ in } \mathbb{E}^d$
 $\Sigma(X) = \lambda \Sigma_1(X) + (1-\lambda) \Sigma_2(X)\}$
- When mapping F by ϕ , the image of F is the line $\phi(F)$ that connects $\phi(\Sigma_1)$ to $\phi(\Sigma_2)$



Types of Pencils




Concentric pencil



Pencil with two limit points



Tangent pencil



Pencil with supporting sphere

Types of pencils of spheres

- If $\phi(F)$ intersects P at 1 point, F contains a single sphere of radius 0, F is a pencil of concentric spheres
- If 2 intersection points, 2 limit points of radius 0
- If 1 tangent point, as if 2 limit points touching, a tangent pencil

Types of pencils


- If $\phi(F)$ does not intersect P , exists some family of hyperplanes tangent to P containing F .
- $\phi(\Sigma_F)$ is the set of tangent points
- Σ_F is the set of points that belong to all spheres in F
- All d -spheres in F intersect on the $(d-1)$ -sphere Σ_1 intersect Σ_2
- Σ_F is the supporting sphere of the pencil F

Radical Hyperplane

- Any point in H_{12} of Σ_1 and Σ_2 has same power with respect to any particular sphere
- Therefore, H_{12} is the radical hyperplane of a pencil of spheres, or of any 2 spheres in the pencil


Radical Hyperplane of types of pencils

- Concentric pencil- does not exist
- Pencil with 2 limit points- the bisector
- Tangent pencil- hyperplane bisector
- Supporting Sphere- the affine hull of the supporting sphere



VD's in Hyperbolic Spaces

- The Poincare model of hyperbolic space:
 - $\mathbb{H}^d = \{X \text{ in } \mathbb{E}^d : x_d > 0\}$
 - A half-space
- Hyperbolic distance: sufficient to decide whether B or C closer to A




Hyperbolic Distance

- Consider a pencil F_A , with limit points A and A' , where A' denotes the symmetric of A w.r.t. hyperplane H_0 of equation $x_d=0$, the radical hyperplane of F_A
- Distance is the radius of sphere of F_A passing through the point.

Hyperbolic VD

- Given n sites in the poincare half-space \mathbb{H}^d , corresponds a region

$$V_h(M_i) = \{X \text{ in } \mathbb{H}^d, \delta_h(X, M_i) \leq \delta_h(X, M_j) \text{ for any } j \neq i\}$$




Projecting the Hyperbolic Diagram

- Identify with half-hyperplane $x_{d+1}=0, x_d > 0$.
The hyperplane H_0 is now identified with the subspace $\{x_{d+1}=x_d=0\}$.
- F_X is mapped to a line in \mathbb{E}^{d+1} parallel to the x_d axis.
- X, X' symmetric w.r.t. H_0 , F_X has limit points at X, X' mapped to $\phi(X), \phi(X')$, symmetric to hyperplane $x_d=0$ in \mathbb{E}^{d+1}

Hyperbolic Diagram in \mathbb{E}^{d+1}

- A point X belongs to $V_h(M_i)$ iff the ray parallel to the x_d axis in \mathbb{E}^{d+1} originating at $\phi(X)$ (entirely contained in P), directed toward $x_d > 0$ intersects the hyperplane $\phi(M_i)^*$ before any other $\phi(M_j)^*$

Hyperbolic Diagram in \mathbb{E}^{d+1}



Consequences of the Projection

- Bisecting surface of 2 points for hyperbolic distance is a half-sphere
 - X is only equidistant between A and B if F_X contains a sphere passing through A and B
 - $\phi(F_X)$ intersects the intersection of $\phi(A)^*$ and $\phi(B)^*$
 - We say that $\phi(X)$ intersects Γ , the projection of the intersection of $\phi(A)^*$ and $\phi(B)^*$ parallel to x_d axis onto P

Bisecting surface

- Γ , the projection of the intersection of $\phi(A)^*$ and $\phi(B)^*$ parallel to x_d axis onto P
- Also the intersection of a hyperplane parallel to x_d and P , which projects to \mathbb{E}^d as a sphere Σ_{AB} , belong to pencil with limit points A, B
- Spheres on that pencil are mapped to $\phi(\Sigma) = \lambda\phi(A) + (1-\lambda)\phi(B)$, with corresponding polar hyperplanes, which all contain the intersection of $\phi(A)^*$ and $\phi(B)^*$

Bisecting Surface

- H is therefore a hyperplane polar to a sphere in F_{AB} , limit points AB .
- H has a hyperplane polar to $\phi(\Sigma_{AB})$, by 17.2.2, so Σ_{AB} belongs to F_{AB}
- Therefore, Σ_{AB} is the sphere in F_{AB} centered on H_0 (unique)

Consequences (2)

- X is equidistant from $d+1$ points, A_0, \dots, A_d , iff $\phi(X)$ is the projection of the intersection of their projected polar halfplanes (intersection $\phi(A_i)^*$) parallel to x_d axis, onto P
- The point at equal hyperbolic distance from $d+1$ points is the limit point of the pencil containing the sphere circumscribed to the $d+1$ points of radical hyperplane H_0

Consequences (3)

- HD can be obtained by projecting $V(M) = \text{intersection } \phi(A_i)^{*+}$ parallel to x_d axis onto P , then vertically onto $x_{d+1}=0$
- This double projection creates injection between VD and HD
- Can do in 1 projection
- Complexity is $O(n^{\text{ceil}(d/2)})$, computed in $O(n \log n + n^{\text{ceil}(d/2)})$ time