

Delaunay Triangulations

Delaunay Triangulations

Guest Lecture by Prof. Andrew P. Black

based on slides by Marc van Kreveld
Utrecht University

Reference

Mark de Berg, Otfried Cheong, Marc van Kreveld, Mark Overmars


Computational Geometry: Algorithms and Applications

Third Edition
Springer, 2008

Chapter 9

Motivation: Terrains by interpolation


To build a model of the terrain surface, we can start with a number of sample points where we know the height.



Motivation: Terrains

How do we interpolate the height at other points?

- Nearest neighbor interpolation
- Piecewise linear interpolation by a triangulation
- Moving windows interpolation
- Natural neighbor interpolation
- ...




Triangulation

Let $P = \{p_1, \dots, p_n\}$ be a point set.
A **triangulation** of P is a maximal
planar subdivision with vertex set P .

Complexity:


- $2n - 2 - k$ triangles
- $3n - 3 - k$ edges

where k is the number of points in P
on the convex hull of P





Triangulation

But which triangulation?



Triangulation




Triangulation


For interpolation, it is good if triangles are not long and skinny. We want to use triangles with *large angles* in our triangulation.

Angle Vector of a Triangulation

- Let \mathcal{T} be a triangulation of P with m triangles. Its **angle vector** is $A(\mathcal{T}) = (\alpha_1, \dots, \alpha_{3m})$ where $\alpha_1, \dots, \alpha_{3m}$ are the angles of \mathcal{T} sorted by increasing value.
- Let \mathcal{T}' be another triangulation of P . We define $A(\mathcal{T}) > A(\mathcal{T}')$ if $A(\mathcal{T})$ is lexicographically larger than $A(\mathcal{T}')$
- \mathcal{T} is **angle optimal** if $A(\mathcal{T}) \geq A(\mathcal{T}')$ for all triangulations \mathcal{T}' of P



Edge Flipping




- Change in angle vector:
 $\alpha_1, \dots, \alpha_6$ are replaced by $\alpha'_1, \dots, \alpha'_6$
- The edge $e = \overline{p_i p_j}$ is **illegal** if $\min_{1 \leq i \leq 6} \alpha_i < \min_{1 \leq i \leq 6} \alpha'_i$
- Flipping an illegal edge increases the angle vector

Characterisation of Illegal Edges

How do we determine if an edge is illegal?

Lemma: The edge $\overline{p_i p_j}$ is illegal if and only if p_l lies in the interior of the circle C .




Inscribed Angle Theorem

Theorem: Let C be a circle, ℓ a line intersecting C in points a and b , and p, q, r, s points lying on the same side of ℓ . Suppose that p, q lie on C , r lies inside C , and s lies outside C . Then

$$\angle arb > \angle apb = \angle aqb > \angle asb,$$

where $\angle abc$ denotes the (smaller) angle defined by the points a, b, c .



Legal Triangulations

A **legal triangulation** is a triangulation that does not contain any illegal edge.

Algorithm LEGALTRIANGULATION(\mathcal{T})

Input. A triangulation \mathcal{T} of a point set P .

Output. A legal triangulation of P .

1. **while** \mathcal{T} contains an illegal edge $\overline{p_i p_j}$
2. **do** (* Flip $\overline{p_i p_j}$ *)
3. Let $p_i p_j p_k$ and $p_i p_j p_l$ be the two triangles adjacent to $\overline{p_i p_j}$.
4. Remove $\overline{p_i p_j}$ from \mathcal{T} , and add $\overline{p_k p_l}$ instead.
5. **return** \mathcal{T}

Question: Why does this algorithm terminate?

“Too slow to be interesting”


Voronoi Diagram and Delaunay Graph

Let P be a set of n points in the plane

The **Voronoi diagram** $\text{Vor}(P)$ is the subdivision of the plane into Voronoi cells $\mathcal{V}(p)$ for all $p \in P$

Let \mathcal{G} be the *dual graph* of $\text{Vor}(P)$

The **Delaunay graph** $\mathcal{DG}(P)$ is the *straight line embedding* of \mathcal{G}




Voronoi Diagram and Delaunay Graph

Let P be a set of n points in the plane

The **Voronoi diagram** $\text{Vor}(P)$ is the subdivision of the plane into Voronoi cells $\mathcal{V}(p)$ for all $p \in P$

Let \mathcal{G} be the *dual graph* of $\text{Vor}(P)$

The **Delaunay graph** $\mathcal{DG}(P)$ is the *straight line embedding* of \mathcal{G}




Voronoi Diagram and Delaunay Graph

Let P be a set of n points in the plane

The **Voronoi diagram** $\text{Vor}(P)$ is the subdivision of the plane into Voronoi cells $\mathcal{V}(p)$ for all $p \in P$

Let \mathcal{G} be the *dual graph* of $\text{Vor}(P)$

The **Delaunay graph** $\mathcal{DG}(P)$ is the *straight line embedding* of \mathcal{G}




Voronoi Diagram and Delaunay Graph

Let P be a set of n points in the plane

The **Voronoi diagram** $\text{Vor}(P)$ is the subdivision of the plane into Voronoi cells $\mathcal{V}(p)$ for all $p \in P$


Let \mathcal{G} be the *dual graph* of $\text{Vor}(P)$

The **Delaunay graph** $\mathcal{DG}(P)$ is the *straight line embedding* of \mathcal{G}




Planarity of the Delaunay Graph

Theorem: The Delaunay graph of a planar point set is a plane graph.




Delaunay Triangulation

If the point set P is in *general position* then the Delaunay graph is a triangulation.




Empty Circle Property

Theorem: Let P be a set of points in the plane, and let \mathcal{T} be a triangulation of P . Then \mathcal{T} is a Delaunay triangulation of P **if and only if** the circumcircle of any triangle of \mathcal{T} does not contain a point of P in its interior.



Delaunay Triangulations and Legal Triangulations

Theorem: Let P be a set of points in the plane. A triangulation \mathcal{T} of P is legal **if and only if** \mathcal{T} is a Delaunay triangulation.



Angle Optimality and Delaunay Triangulations

Theorem: Let P be a set of points in the plane. Any angle-optimal triangulation of P is a Delaunay triangulation of P . Furthermore, any Delaunay triangulation of P maximizes the minimum angle over all triangulations of P .

Computing Delaunay Triangulations

There are several ways to compute the Delaunay triangulation:

- By iterative flipping from any triangulation
- By plane sweep
- By randomized incremental construction
- By conversion from the Voronoi diagram

The last three run in $O(n \log n)$ time [expected] for n points in the plane


Randomized Incremental Alg. for Delaunay Triangulation

- start with a large bounding triangle
- add points in random order, maintaining the Delaunay triangulation
- remove bounding points

Construct Bounding Triangle

Bounding triangle avoids inconvenient edge cases.


Add leftmost point p_{-2} and rightmost point p_{-1} ; form a triangle with top point p_0



Add Points in Random Order

Add a point p_r :


- 1 find triangle $p_i p_j p_k$ containing p_r
- 2 add new edges from p_i to p_r , p_j to p_r , p_k to p_r
- 3 fix-up the new triangles if they are too skinny



Add Points in Random Order

Add a point p_r :


- 1 p_r falls on an existing edge $e = p_i p_j p_k$
- 2 add new edges from p_r to opposite vertices in *both* triangles sharing e
- 3 fix-up the new triangles if they are too skinny




Fix-up the new triangles

Find illegal edges and flip them:

- Which edges can become illegal due to insertion of p_r ?
 - Only those edges adjacent to new triangles
- $\text{LEGALIZEEDGE}(p_r, \overline{p_i p_j}, \mathcal{T})$ recursively makes all edges affected by p_r in \mathcal{T} legal.




Fix-up the new triangles



LEGALIZEEDGE

 $\text{LEGALIZEEDGE}(p_r, \overline{p_i p_j}, \mathcal{T})$

1. (* The point being inserted is p_r , and $\overline{p_i p_j}$ is the edge of \mathcal{T} that may need to be flipped. *)
2. **if** $\overline{p_i p_j}$ is illegal
3. **then** Let $p_i p_j p_k$ be the triangle adjacent to $p_r p_i p_j$ along $\overline{p_i p_j}$.
4. (* Flip $\overline{p_i p_j}$: *) Replace $\overline{p_i p_j}$ with $\overline{p_r p_k}$.
5. $\text{LEGALIZEEDGE}(p_r, \overline{p_i p_k}, \mathcal{T})$
6. $\text{LEGALIZEEDGE}(p_r, \overline{p_k p_j}, \mathcal{T})$

How to test if $\overline{p_i p_j}$ is legal?

Algorithm DELAUNAYTRIANGULATION(P)

Input. A set P of $n + 1$ points in the plane.


Output. A Delaunay triangulation of P .

1. Let p_0 be the lexicographically highest point of P , that is, the rightmost among the points with largest y -coordinate.
2. Let p_{-1} and p_{-2} be two points in \mathbb{R}^2 sufficiently far away and such that P is contained in the triangle $p_0p_{-1}p_{-2}$.
3. Initialize \mathcal{T} as the triangulation consisting of the single triangle $p_0p_{-1}p_{-2}$.
4. Compute a random permutation p_1, p_2, \dots, p_n of $P \setminus \{p_0\}$.
5. **for** $r \leftarrow 1$ **to** n
6. **do** (* Insert p_r into \mathcal{T} : *)
7. Find a triangle $p_i p_j p_k \in \mathcal{T}$ containing p_r .
8. **if** p_r lies in the interior of the triangle $p_i p_j p_k$
9. **then** Add edges from p_r to the three vertices of $p_i p_j p_k$, thereby splitting $p_i p_j p_k$ into three triangles.
10. LEGALIZEEDGE($p_r, \overline{p_i p_j}, \mathcal{T}$)
11. LEGALIZEEDGE($p_r, \overline{p_j p_k}, \mathcal{T}$)
12. LEGALIZEEDGE($p_r, \overline{p_i p_k}, \mathcal{T}$)


6. **do** (* Insert p_r into \mathcal{T} : *)
7. Find a triangle $p_i p_j p_k \in \mathcal{T}$ containing p_r .
8. **if** p_r lies in the interior of the triangle $p_i p_j p_k$
9. **then** Add edges from p_r to the three vertices of $p_i p_j p_k$, thereby splitting $p_i p_j p_k$ into three triangles.
10. LEGALIZEEDGE($p_r, \overline{p_i p_j}$, \mathcal{T})
11. LEGALIZEEDGE($p_r, \overline{p_j p_k}$, \mathcal{T})
12. LEGALIZEEDGE($p_r, \overline{p_k p_i}$, \mathcal{T})
13. **else** (* p_r lies on an edge of $p_i p_j p_k$, say the edge $\overline{p_i p_j}$ *)
14. Add edges from p_r to p_k and to the third vertex p_l of the other triangle that is incident to $\overline{p_i p_j}$, thereby splitting the two triangles incident to $\overline{p_i p_j}$ into four triangles.
15. LEGALIZEEDGE($p_r, \overline{p_i p_l}$, \mathcal{T})
16. LEGALIZEEDGE($p_r, \overline{p_l p_j}$, \mathcal{T})
17. LEGALIZEEDGE($p_r, \overline{p_j p_k}$, \mathcal{T})
18. LEGALIZEEDGE($p_r, \overline{p_k p_i}$, \mathcal{T})
19. Discard p_{-1} and p_{-2} with all their incident edges from \mathcal{T} .
20. **return** \mathcal{T}

How does one “find a triangle”?


Build a Point Location Graph



Build a Point Location Graph



Build a Point Location Graph




Using Delaunay Triangulations

Delaunay triangulations help in constructing various things:

- Euclidean Minimum Spanning Trees
- Approximations to the Euclidean Traveling Salesperson Problem
- α -Hulls


Euclidean Minimum Spanning Tree

For a set P of n points in the plane, the **Euclidean Minimum Spanning Tree** is the graph with minimum summed edge length that connects all points in P and has only the points of P as vertices



Euclidean Minimum Spanning Tree

For a set P of n points in the plane, the **Euclidean Minimum Spanning Tree** is the graph with minimum summed edge length that connects all points in P and has only the points of P as vertices



Euclidean Minimum Spanning Tree

Lemma: The Euclidean Minimum Spanning Tree does not have cycles (it really is a tree)


Proof: Suppose G is the shortest connected graph and it has a cycle. Removing one edge from the cycle makes a new graph G' that is still connected but which is shorter. Contradiction

Euclidean Minimum Spanning Tree

Lemma: Every edge of the Euclidean Minimum Spanning Tree is an edge in the Delaunay graph

Proof: Suppose T is an EMST with an edge $e = \overline{pq}$ that is not Delaunay

Consider the circle C that has e as its diameter. Since e is not Delaunay, C must contain another point r in P (different from p and q)




Euclidean Minimum Spanning Tree

Lemma: Every edge of the Euclidean Minimum Spanning Tree is an edge in the Delaunay graph

Proof: (continued)

Either the path in T from r to p passes through q , or vice versa. The cases are symmetric, so we can assume the former case




Euclidean Minimum Spanning Tree

Lemma: Every edge of the Euclidean Minimum Spanning Tree is an edge in the Delaunay graph

Proof: (continued)

Then removing e and inserting \overline{pr} instead will give a connected graph again (in fact, a tree)

Since q was the furthest point from p inside C , r is closer to q , so T was not a *minimum* spanning tree.
Contradiction



Euclidean Minimum Spanning Tree

How can we compute a Euclidean Minimum Spanning Tree efficiently?

Union-Find data structure: maintains disjoint sets and allows these operations:

- **Union**: Takes two sets and makes one new set that is the union (destroys the two given sets)
- **Find**: Takes one element and returns the name of the set that contains it

If there are n elements in total, then all **Unions** together take $O(n \log n)$ time and each **Find** operation takes $O(1)$ time

Euclidean Minimum Spanning Tree

Let P be a set of n points in the plane for which we want to compute the EMST

- 1 Make a Union-Find structure where every point of P is in a separate set
- 2 Construct the Delaunay triangulation DT of P
- 3 Take all edges of DT and sort them by length
- 4 For all edges e from short to long:
 - Let the endpoints of e be p and q
 - If $\mathbf{Find}(p) \neq \mathbf{Find}(q)$, then put e in the EMST, and $\mathbf{Union}(\mathbf{Find}(p), \mathbf{Find}(q))$

Euclidean Minimum Spanning Tree


Step 1 takes linear time, the other three steps take $O(n \log n)$ time

Theorem: Let P be a set of n points in the plane. The Euclidean Minimum Spanning Tree of P can be computed in $O(n \log n)$ time

The traveling salesperson problem

Given a set P of n points in the plane, the **Euclidean Traveling Salesperson Problem** is to compute a tour (cycle) that visits all points of P and has minimum length

A tour is an *order* on the points of P (more precisely: a cyclic order). A set of n points has $(n - 1)!$ different tours



The traveling salesperson problem

We can determine the length of each tour in $O(n)$ time: a brute-force algorithm to solve the Euclidean Traveling Salesperson Problem (ETSP) takes $O(n) \cdot O((n-1)!) = O(n!)$ time

How bad is $n!$?

Efficiency

n	n^2	2^n	$n!$
6	36	64	720
7	49	128	5040
8	64	256	40K
9	81	512	360K
10	100	1024	3.5M
15	225	32K	2,000,000T
20	400	1M	
30	900	1G	

Clever algorithms can solve instances in $O(n^2 \cdot 2^n)$ time

Approximation algorithms

If an algorithm A solves an optimization problem always within a factor k of the optimum, then A is called an **k -approximation algorithm**

If an instance I of ETSP has an optimal solution of length L , then a k -approximation algorithm will find a tour of length $\leq k \cdot L$

Approximation algorithms

Consider the diameter problem of a set of n points. We can compute the real value of the diameter in $O(n \log n)$ time

Suppose we take any point p , determine its furthest point q , and return their distance. This takes only $O(n)$ time


Question: Is this an approximation algorithm?

Approximation algorithms

Consider the diameter problem of a set of n points. We can compute the real value of the diameter in $O(n \log n)$ time

Suppose we take any point p , determine its furthest point q , and return their distance. This takes only $O(n)$ time

Question: Is this an approximation algorithm?



Approximation algorithms


Suppose we determine the point with minimum x -coordinate p and the point with maximum x -coordinate q , and return their distance. This takes only $O(n)$ time

Question: Is this an approximation algorithm?

Approximation algorithms

Suppose we determine the point with minimum x -coordinate p and the point with maximum x -coordinate q , and return their distance. This takes only $O(n)$ time

Question: Is this an approximation algorithm?



Approximation algorithms

Suppose we determine the point with minimum x -coordinate p and the point with maximum x -coordinate q .

Then we determine the point with minimum y -coordinate r and the point with maximum y -coordinate s .

We return $\max(d(p,q), d(r,s))$.

This takes only $O(n)$ time

Question: Is this an approximation algorithm?

Approximation algorithms


Suppose we determine the point with minimum x -coordinate p and the point with maximum x -coordinate q .

Then we determine the point with minimum y -coordinate r and the point with maximum y -coordinate s .

We return $\max(d(p,q), d(r,s))$.

This takes only $O(n)$ time


Question: Is this an approximation algorithm?



Approximation algorithms

Back to Euclidean Traveling Salesperson:

We will use the EMST to approximate the ETSP




start at any vertex

Approximation algorithms

Back to Euclidean Traveling Salesperson:

We will use the EMST to approximate the ETSP




follow an edge on one side

Approximation algorithms

Back to Euclidean Traveling Salesperson:

We will use the EMST to approximate the ETSP




... to get to another vertex

Approximation algorithms

Back to Euclidean Traveling Salesperson:

We will use the EMST to approximate the ETSP




proceed this way

Approximation algorithms

Back to Euclidean Traveling Salesperson:

We will use the EMST to approximate the ETSP




proceed this way

Approximation algorithms

Back to Euclidean Traveling Salesperson:

We will use the EMST to approximate the ETSP




proceed this way

Approximation algorithms

Back to Euclidean Traveling Salesperson:

We will use the EMST to approximate the ETSP




skipping visited vertices

Approximation algorithms

Back to Euclidean Traveling Salesperson:

We will use the EMST to approximate the ETSP




skipping visited vertices

Approximation algorithms

Back to Euclidean Traveling Salesperson:

We will use the EMST to approximate the ETSP




skipping visited vertices

Approximation algorithms

Back to Euclidean Traveling Salesperson:

We will use the EMST to approximate the ETSP




skipping visited vertices

Approximation algorithms

Back to Euclidean Traveling Salesperson:

We will use the EMST to approximate the ETSP




skipping visited vertices

Approximation algorithms

Back to Euclidean Traveling Salesperson:

We will use the EMST to approximate the ETSP




skipping visited vertices

Approximation algorithms

Back to Euclidean Traveling Salesperson:

We will use the EMST to approximate the ETSP




skipping visited vertices

Approximation algorithms

Back to Euclidean Traveling Salesperson:

We will use the EMST to approximate the ETSP




skipping visited vertices

Approximation algorithms

Back to Euclidean Traveling Salesperson:

We will use the EMST to approximate the ETSP




and close the tour

Approximation algorithms

Back to Euclidean Traveling Salesperson:

We will use the EMST to approximate the ETSP




and close the tour

Approximation algorithms

Why is this tour an approximation?

- The walk visits every edge twice, so it has length $2 \cdot |EMST|$
- The tour skips vertices, which means the tour has length $\leq 2 \cdot |EMST|$
- The optimal ETSP-tour is a spanning tree if you remove any edge!!!
So $|EMST| < |ETSP|$



Approximation algorithms


Theorem: Given a set of n points in the plane, a tour visiting all points whose length is at most twice the minimum possible can be computed in $O(n \log n)$ time

In other words: an $O(n \log n)$ time, 2-approximation for ETSP exists

α -Shapes

Suppose that you have a set of points in the plane that were sampled from a shape


We would like to reconstruct the shape



α -Shapes

Suppose that you have a set of points in the plane that were sampled from a shape

We would like to reconstruct the shape




α -Shapes

An α -disk is a disk of radius α

The α -shape of a point set P is the graph with:

- the points of P as the vertices, and
- vertices p, q are connected by an edge if there exists an α -disk with p and q on the disk's boundary but no other points of P inside, or on the boundary, of the disk.




α -Shapes

An α -disk is a disk of radius α

The α -shape of a point set P is the graph with:

- the points of P as the vertices, and
- vertices p, q are connected by an edge if there exists an α -disk with p and q on the disk's boundary but no other points of P inside, or on the boundary, of the disk.




α -Shapes

An α -disk is a disk of radius α

The α -shape of a point set P is the graph with:

- the points of P as the vertices, and
- vertices p, q are connected by an edge if there exists an α -disk with p and q on the disk's boundary but no other points of P inside, or on the boundary, of the disk.




α -Shapes

An α -disk is a disk of radius α

The α -shape of a point set P is the graph with:


- the points of P as the vertices, and
- vertices p, q are connected by an edge if there exists an α -disk with p and q on the disk's boundary but no other points of P inside, or on the boundary, of the disk.



α -Shapes


Because of the empty disk property of Delaunay triangulations (each Delaunay edge has an empty disk through its endpoints), every α -shape edge is also a Delaunay edge

Hence: there are $O(n)$ α -shape edges, and they cannot properly intersect




α -Shapes

Given the Delaunay triangulation, we can determine for any edge all sizes of empty disks through the endpoints in $O(1)$ time



So the α -shape can be computed in $O(n \log n)$ time



Conclusions

The **Delaunay triangulation** is a versatile structure that can be computed in $O(n \log n)$ time for a set of n points in the plane.

Approximation algorithms are like heuristics, but they come with a guarantee on the quality of the approximation. They are useful when an computing an optimal solution is too time-consuming.