

Preferred directions for resolving the non-uniqueness of Delaunay triangulations

Christopher Dyken and Michael S. Floater

Abstract: This note proposes a simple rule to determine a unique triangulation among all Delaunay triangulations of a planar point set, based on two *preferred directions*. We show that the triangulation can be generated by extending Lawson’s edge-swapping algorithm and that point deletion is a local procedure. The rule can be implemented exactly when the points have integer coordinates and can be used to improve image compression methods.

Key words: Delaunay triangulation, uniqueness, edge-swapping, integer arithmetic.

1. Introduction

Delaunay triangulations [3] play an important role in computational geometry [9]. A recent application that has emerged is their use in compressing digital images [4]. Such images are represented by rectangular arrays of grey scale values or colour values and one approach to compression is to start by representing them as piecewise linear functions over regular triangulations and then to approximate these functions by piecewise linear functions over triangulations of subsets of the points. If one could agree on a *unique* method of triangulating the points, one would obtain higher compression rates because the sender would only need to encode the points and the height values, not the connectivity of the triangulation: the receiver would be able to reproduce the triangulation exactly.

One advantage of Delaunay triangulations is that they are *almost* unique. In fact they are unique for point sets containing no sets of four co-circular points. However, in the case that a set of planar points is a subset of a rectangular array of points, there will typically be many co-circular points and therefore a large number of Delaunay triangulations.

An obvious and simple approach to the non-uniqueness problem is to perturb the points randomly before triangulating but there is no guarantee that the perturbed points will have a unique Delaunay triangulation. An alternative approach could be a symbolic perturbation method as discussed in [1,5,6,8,10]. Such a method should lead to a unique Delaunay triangulation of the ‘perturbed’ points, but may not be a valid triangulation of the original ones.

The purpose of this note is to point out that non-uniqueness can be resolved without perturbing the points. We show that a simple rule based on two preferred directions can be used to determine a *unique* member of all the Delaunay triangulations of a set of points in the plane. We show that the rule can simply be incorporated into Lawson’s swapping algorithm and that point deletion is a local procedure. We further show that, importantly, the rule can be computed exactly in integer arithmetic. The preferred direction method could immediately be applied to the compression algorithms described above.

2. Triangulating quadrilaterals

Let Q be a quadrilateral in the plane, with ordered vertices v_1, w_1, v_2, w_2 as in Figure 1. If Q is strictly convex, i.e., convex and such that no three vertices are collinear, then there




Figure 1: A strictly convex quadrilateral

are two ways to triangulate it, either by placing a diagonal edge between v_1 and v_2 , as in Figure 1, or between w_1 and w_2 . We want to propose a simple rule which determines the diagonal of any such quadrilateral uniquely. A natural rule seems to be to use a *preferred direction*: we choose that diagonal which makes the smallest angle with some arbitrary, fixed straight line. However, the two angles could be equal, and in order to distinguish the two diagonals in this case, we will use a second preferred direction. Thus we choose any two non-zero vectors d_1 and d_2 which are neither parallel nor orthogonal to each other. One such choice would be $d_1 = (1, 0)$ and $d_2 = (1, 1)$. For any line (or line segment) ℓ , let $\alpha_i(\ell)$, where $0 \leq \alpha_i(\ell) \leq \pi/2$, be the angle between ℓ and the (undirected) vector d_i , $i = 1, 2$. Then we define the *score* of ℓ as the ordered pair of angles

$$\text{score}(\ell) := (\alpha_1(\ell), \alpha_2(\ell)).$$

We compare scores lexicographically. Thus for two arbitrary lines ℓ and m , we say that $\text{score}(\ell) < \text{score}(m)$ if either $\alpha_1(\ell) < \alpha_1(m)$ or $\alpha_1(\ell) = \alpha_1(m)$ and $\alpha_2(\ell) < \alpha_2(m)$.

Lemma 2.1. *Two lines ℓ and m have the same score if and only if they are parallel.*

Proof: If ℓ and m are parallel, then clearly $\alpha_i(\ell) = \alpha_i(m)$ for both $i = 1, 2$ and so ℓ and m have the same score. Conversely, suppose ℓ and m are not parallel but that they have the same score. Then at the point p of intersection between ℓ and m , the two lines through p in the directions d_1 and d_2 must bisect the lines ℓ and m . But this can only occur if d_1 is either parallel or orthogonal to d_2 , which is a contradiction. ■

Our *preferred direction rule* for the quadrilateral Q in Figure 1 simply chooses that diagonal, $[v_1, v_2]$ or $[w_1, w_2]$, with the lowest score. Since the two diagonals are never parallel, they always have distinct scores by Lemma 2.1.

3. Triangulating convex polygons

We next use the preferred direction rule to triangulate uniquely any strictly convex polygon P , illustrated in Figure 2. Consider $E_I(P)$, the set of all interior edges of P , i.e., all line segments $[v_1, v_2]$ connecting non-neighbouring pairs of vertices v_1 and v_2 of P . We start by ranking all the edges of $E_I(P)$ according to their score. We then employ an insertion algorithm, inserting edges of $E_I(P)$ into P in order of their ranking. In the first step, we insert all edges of $E_I(P)$ (one or more) which share the lowest score. Note that if there are more than one of these, they must all be parallel by Lemma 2.1, and so they do not cross each other. In the general step, we insert all edges of $E_I(P)$ with the current lowest score




Figure 2: A strictly convex polygon and triangulation

which do not cross edges previously inserted. We continue until we have triangulated P , denoting the triangulation by $\mathcal{T}_{d_1, d_2}(P)$.

We will show that $\mathcal{T}_{d_1, d_2}(P)$ has the useful property that it can be generated by a local optimisation procedure based on edge swapping. Suppose $\mathcal{T}(P)$ is any triangulation of P . Every interior edge e of $\mathcal{T}(P)$ is the diagonal of a strictly convex quadrilateral Q , and can therefore be swapped by the other diagonal e' of Q to form a new triangulation. We say that e is locally optimal if its score is lower than that of e' . Otherwise, we optimize Q by swapping e with e' . We keep applying this local optimisation procedure until no more swaps can be performed, in which case we say that the triangulation is *locally optimal*.

Lemma 3.1. *Let $\mathcal{T}(P)$ be any triangulation of P . If there is an edge $e \in E_I(P)$ which crosses an edge of $\mathcal{T}(P)$ with a higher score, then $\mathcal{T}(P)$ is not locally optimal.*

Proof: Let $e' = [v_1, v_2]$ be an edge of $\mathcal{T}(P)$ with the highest score among all edges of $\mathcal{T}(P)$ which cross e . We will show that e' is not locally optimal. Let p be the intersection of e and e' . Without loss of generality, and by translating and rotating P and d_1 and d_2 about p , we may assume that $p = (0, 0)$ and $d_1 = (1, 0)$. Then since e' has a higher score than e , e' cannot lie along the x -axis, and so we can assume that $y_1 < 0 < y_2$ (where $v_i = (x_i, y_i)$). Further, by reflecting all points, edges, and d_1 and d_2 about the y -axis if necessary, we may assume that $x_1 \leq 0 \leq x_2$; see Figure 3.




Figure 3: The edges e and e'

Next let ℓ' be the infinite straight line passing through e' and let ℓ'' be its reflection about the y -axis, and let A_1 and A_2 be the two open semi-infinite cones bounded by ℓ' and ℓ'' , with A_1 in the positive y half-plane and A_2 in the negative y half-plane; see Figure 4. Clearly because $\alpha_1(e) \leq \alpha_1(\ell') = \alpha_1(\ell'')$, e does not intersect $A_1 \cup A_2$. Thus e is contained




Figure 4: Regions of vertices

in the the union of the two closed regions B_1 and B_2 shown in Figure 4. Note, moreover, that e is not contained in the line l' . It may however be contained in the line l'' .

Let Q be the convex quadrilateral of $\mathcal{T}(P)$ with e' as its diagonal. Let the two vertices in Q other than v_1 and v_2 be w_1 and w_2 , with w_1 lying on the side of l' containing the positive x -axis and w_2 lying on the side containing the negative x -axis; see Figure 5 for an example. Next we show that $w_2 \in B_1$. Indeed if $w_2 \in A_1$ then by the convexity of P , the edge $[v_1, w_2]$ of $\mathcal{T}(P)$ intersects e and $[v_1, w_2]$ has a higher score than e' which contradicts the definition of e' . Furthermore, if $w_2 \in l''$ and $e \not\subset l''$, then again the edge $[v_1, w_2]$ would intersect e and have a higher score than e' . Therefore, if $w_2 \in l''$ then also $e \subset l''$. In this case, since no three vertices of P are collinear, we conclude that w_2 is an end point of e .

Similarly, $w_1 \in B_2$, because if $w_1 \in A_2$ then the edge $[v_2, w_1]$ would intersect e and have a higher score than e' which contradicts the definition of e' . A similar argument shows that if $w_1 \in l''$ then $e \subset l''$ and w_1 must be an end point of e .




Figure 5: A possible quadrilateral

To complete the proof, we will show that e' is not locally optimal by showing that $[w_1, w_2]$ has a lower score than e' . First observe from Figure 5 that $\alpha_1([w_1, w_2]) \leq \alpha_1(e')$. Moreover, if either w_2 is in the interior of B_1 or w_1 is in the interior of B_2 , or both, then the inequality is strict and so $\text{score}([w_1, w_2]) < \text{score}(e')$. The only remaining possibility is that both w_1 and w_2 lie on the line l'' , in which case $\alpha_1([w_1, w_2]) = \alpha_1(e')$. We have shown, however, that in this case both w_1 and w_2 are the end points of the edge e , so that

$e = [w_1, w_2]$. Since e has a lower score than e' by assumption, e' is therefore not locally optimal. ■

An immediate consequence of Lemma 3.1 is

Theorem 3.2. *If $\mathcal{T}(P)$ is locally optimal then $\mathcal{T}(P) = \mathcal{T}_{d_1, d_2}(P)$.*

Proof: Lemma 3.1 implies that if e is an interior edge of $\mathcal{T}(P)$ then the only edges of $E_I(P)$ which cross it have a higher score than e . This means that when the score of e is reached in the insertion algorithm, e will be inserted into P . Thus e is in $\mathcal{T}_{d_1, d_2}(P)$. ■

Lemma 3.1 and Theorem 3.2 together clearly imply

Theorem 3.3. *A line segment $e \in E_I(P)$ is an edge of $\mathcal{T}_{d_1, d_2}(P)$ if and only if it is not crossed by an edge in $E_I(P)$ with a lower score.*

4. Delaunay triangulations

Finally we use the preferred direction rule to determine a unique triangulation among all possible Delaunay triangulations of a given set of points in the plane.

Any set of planar points V which are not all collinear admits a unique Delaunay *pretriangulation*, which is a tiling of the points, whose boundary is the convex hull of V [11]. Two points form an edge in the tiling if and only if they are strong neighbours in the Voronoi diagram of V [11, 9]. Two points are Voronoi *neighbours* if their Voronoi tiles intersect. Two such tiles intersect either in a line segment or a point. If the intersection is a line segment the two points are *strong* neighbours, and they are *weak* neighbours otherwise. The vertices of each tile in the Delaunay pretriangulation lie on a circle. Tiles with three vertices are triangles. By triangulating any tile with four or more vertices *arbitrarily*, we convert the Delaunay pretriangulation into a Delaunay triangulation.

Since each tile P of the Delaunay pretriangulation is strictly convex, we can simply triangulate it using $\mathcal{T}_{d_1, d_2}(P)$ and in this way we determine a unique Delaunay triangulation of V , which we will denote by $\mathcal{T}_{d_1, d_2}(V)$.

Similarly to the case of convex polygons, we next show that given an arbitrary triangulation $\mathcal{T}(V)$, we always reach $\mathcal{T}_{d_1, d_2}(V)$ by edge swapping. We simply augment Lawson's local optimisation procedure for Delaunay triangulations [7] with the preferred direction rule as follows. Suppose $e = [v_1, v_2]$ is an interior edge of some triangulation $\mathcal{T}(V)$. If Q , the quadrilateral having e as its diagonal, is not strictly convex we say that e is locally optimal. Otherwise, if $e' = [w_1, w_2]$ denotes the opposite diagonal of Q , we say that e is locally optimal if w_2 lies strictly outside the circumcircle C through v_1, v_2, w_1 . If w_2 lies strictly inside C , we swap e with e' . Otherwise v_1, v_2, w_1, w_2 are co-circular and we use the preferred direction rule as a tie-breaker: e is locally optimal if $\text{score}(e) < \text{score}(e')$. Otherwise we swap e with e' . We say that $\mathcal{T}(V)$ is locally optimal if all its interior edges are locally optimal.

Theorem 4.1. *If $\mathcal{T}(V)$ is locally optimal then $\mathcal{T}(V) = \mathcal{T}_{d_1, d_2}(V)$.*

Proof: If $\mathcal{T}(V)$ is locally optimal then it is a Delaunay triangulation [7], [11]. Thus every edge of $\mathcal{T}(V)$ which is in the Delaunay pretriangulation of V is also in $\mathcal{T}_{d_1, d_2}(V)$.

Every remaining edge e of $\mathcal{T}(V)$ is an interior edge of some tile P of the pretriangulation. But since e is then in $\mathcal{T}_{d_1, d_2}(P)$ it must also be in $\mathcal{T}_{d_1, d_2}(V)$. ■

Note further that an important operation on Delaunay triangulations is point deletion. It is well known that when an interior vertex of a Delaunay triangulation is deleted, a Delaunay triangulation of the remaining points can be constructed by simply retriangulating the hole created by the removal of v , and thus updating the triangulation is a *local* operation, and can be implemented efficiently. We now show that the triangulation $\mathcal{T}_{d_1, d_2}(V)$ has an analogous property.

Theorem 4.2. *Let $v \in V$ be an interior vertex of the triangulation $\mathcal{T}_{d_1, d_2}(V)$. Then every edge e of $\mathcal{T}_{d_1, d_2}(V)$ which is not incident on v also belongs to the triangulation $\mathcal{T}_{d_1, d_2}(V \setminus v)$.*

Proof: Suppose first that e is an edge of the Delaunay pretriangulation of V . Then its end points are strong neighbours in the Voronoi diagram of V , and are therefore also strong neighbours of the Voronoi diagram of $V \setminus v$, and so e is also an edge of the Delaunay pretriangulation of $V \setminus v$, and therefore an edge of $\mathcal{T}_{d_1, d_2}(V \setminus v)$.

The remaining possibility is that e is an interior edge of some strictly convex tile P of the Delaunay pretriangulation of V , and all vertices of P lie on a circle. If v is not a vertex of P , then P will also be a tile in the Delaunay pretriangulation of $V \setminus v$, and so e , being an interior edge of $\mathcal{T}_{d_1, d_2}(P)$, will be contained in $\mathcal{T}_{d_1, d_2}(V \setminus v)$. Otherwise v is a vertex of P . Then the polygon P' , formed by removing v from P , is a tile in the Delaunay pretriangulation of $V \setminus v$. Now if v and the two end points of e are consecutive vertices of P , then e will be an edge of P' , and therefore of $\mathcal{T}_{d_1, d_2}(V \setminus v)$. Otherwise, since $E_I(P') \subset E_I(P)$, we see by Theorem 3.3 that since e is not crossed by any edge in $E_I(P)$ with a lower score, it is not crossed by any edge of $E_I(P')$ with a lower score and is therefore an interior edge of $\mathcal{T}_{d_1, d_2}(P')$, and hence also an edge of $\mathcal{T}_{d_1, d_2}(V \setminus v)$. ■

5. Numerical implementation

Using the implementation of the circumcircle test proposed by Cline and Renka [2], one can construct a Delaunay triangulation of a set V of points with integer coordinates using integer arithmetic. So in order to construct the triangulation $\mathcal{T}_{d_1, d_2}(V)$ in this case, we only need to show that the preferred direction rule can be implemented exactly. To see that this is possible, consider the angle test

$$\alpha_i([v_1, v_2]) < \alpha_i([w_1, w_2]), \quad i = 1, 2.$$

One way to convert this to an integer comparison is to observe that it is equivalent to

$$\cos \alpha_i([v_1, v_2]) > \cos \alpha_i([w_1, w_2])$$

which, using scalar products, is equivalent to

$$\left| \frac{(v_2 - v_1) \cdot d_i}{|v_2 - v_1| \cdot |d_i|} \right| > \left| \frac{(w_2 - w_1) \cdot d_i}{|w_2 - w_1| \cdot |d_i|} \right|.$$

By squaring and removing the common denominator we get the equivalent test

$$|w_2 - w_1|^2 ((v_2 - v_1) \cdot d_i)^2 - |v_2 - v_1|^2 ((w_2 - w_1) \cdot d_i)^2 > 0.$$

Since the left hand side is a polynomial in the coordinates of the points v_1, w_1, v_2, w_2 , and the vector d_i , we see that provided d_1 and d_2 also have integer coordinates, the left hand side is also an integer. Clearly, all angle tests involved in comparing two scores involve testing the sign (positive, zero, or negative) of such an integer.

Figure 6 shows the Delaunay triangulation $\mathcal{T}_{d_1, d_2}(V)$ for two different choices of d_1 and d_2 , where V is a subset of points on a square grid. The triangulations in the figure were found by recursive point insertion, applying Lawson’s swapping algorithm augmented with the preferred direction rule, and using integer arithmetic. One could, however, use any method to find one of the possible Delaunay triangulations and then apply edge swapping with the preferred direction rule to reach the triangulation $\mathcal{T}_{d_1, d_2}(V)$. Since the edges of the Delaunay pretriangulation will never be swapped, the only swapping required will occur inside the tiles of the pretriangulation, and so the number of swaps will generally be low. The Delaunay triangulations computed in [4] were found by recursive point removal and this algorithm could be augmented with the preferred direction rule and the updates would be local due to Theorem 4.2.




Figure 6: Triangulations with $d_1 = (1, 0)$, $d_2 = (1, 1)$ and $d_1 = (1, -1)$, $d_2 = (0, 1)$

References

1. P. Alliez, O. Devillers, and J. Snoeyink, Removing degeneracies by perturbing the problem or the world, *Reliable Computing* **6** (2000), 61–79.
2. A. K. Cline and R. J. Renka, A storage efficient method for construction of a Thiessen triangulation, *Rocky Mountain J. Math*, **14** (1984), 119–140.
3. B. Delaunay, Sur la sphère vide, *Bull. Acad. Sci. USSR (VII)*, *Classe Sci. Mat. Nat.*, (1934), 793–800.
4. L. Demaret, N. Dyn, M. S. Floater and A. Iske, Adaptive Thinning for Terrain Modelling and Image Compression, in *Advances in Multiresolution for Geometric Modelling*, N. A. Dodgson, M. S. Floater, and M. A. Sabin (eds.), Springer, Heidelberg, 2004, 319–338.

5. O. Devillers and M. Teillaud, Perturbations and vertex removal in a 3D delaunay triangulation, in SODA 2003, pp. 313–319.
6. H. Edelsbrunner and E. P. Mücke, Simulation of simplicity: a technique to cope with degenerate cases in geometric algorithms, *ACM Trans. Graphics* **9** (1990), 67–104.
7. C. F. Lawson, Software for C^1 surface interpolation, in *Mathematical Software III*, J. Rice (ed.), Academic Press, 1977, pp. 161–194.
8. E. P. Mücke, A robust implementation for three-dimensional Delaunay triangulations, *Int. J. Comput. Geom. Appl.* **8** (1998), 255–276.
9. J. O’Rourke, *Computational Geometry in C (Second Edition)*, Cambridge University Press, 1998.
10. R. Seidel, The nature and meaning of perturbations in geometric computing, *Discrete Comput. Geom.* **19** (1998), 1–17.
11. R. Sibson, Locally equiangular triangulations, *Comput. J.* **21** (1978), 243–245.

Christopher Dyken and Michael S. Floater
Centre of Mathematics for Applications
Department of Informatics
University of Oslo
Postbox 1053, Blindern
0316 Oslo, NORWAY
dyken@cma.uio.no
michaelf@ifi.uio.no