

Delaunay Triangulations and Voronoi Diagrams for Riemannian Manifolds

Greg Leibon
Department of Mathematics
Dartmouth College
Hanover, NH 03755
leibon@math.dartmouth.edu

David Letscher
Department of Mathematics
Oklahoma State University
Stillwater, OK 74078
letscher@math.okstate.edu

ABSTRACT

For a sufficiently dense set of points in any closed Riemannian manifold, we prove that a unique Delaunay triangulation exists. This triangulation has the same properties as in Euclidean space. Algorithms for constructing these triangulations will also be described.

1. INTRODUCTION

Given a set of points in Euclidean space, Delaunay triangulations give a canonical triangulation whose vertices are these points. These triangulations have many nice combinatorial and geometric properties that make them extremely useful. These triangulations can also be constructed in Riemannian manifolds. However, they do not exist for arbitrary set of points; there are certain density requirements to ensure that the triangulation can accurately represent both the topology and geometry of the manifold. These triangulations are canonically determined by the set of points and have many of the properties that Delaunay triangulations have in the Euclidean space.

A familiarity with Delaunay triangulations in Euclidean space will be assumed. There are several reasons that working in Riemannian geometry makes these problems more difficult. One of the the key properties of Delaunay triangulations is the empty circumscribing sphere condition. In \mathbf{R}^n there is a unique sphere with any given $n + 1$ points on its boundary. However, in Riemannian manifolds, there can be many spheres through $n + 1$ points and they are not necessarily embedded. Also, measuring distances can be more difficult since there are no unique "lines" connecting any pair of points. These difficulties and others can be overcome if all work is done in sufficiently small neighborhoods where the differences from Euclidean space can be minimized. If we do not make these extra assumptions we can still obtain well defined structures on M ; however, we cannot ensure that it

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Computational Geometry 2000 Hong Kong China
Copyright ACM 2000 1-58113-224-7/00/6...\$5.00

will be a triangulation and that it will be Delaunay.

First, we will discuss the necessary portions of Riemannian geometry that will be needed to prove the existence of Delaunay triangulations. Then we will introduce the density requirements for Delaunay triangulations to exist and be well defined. The next step will be to show that Voronoi diagrams will act as they do in Euclidean space with the same density conditions. We will also discuss how these theorems can be modified to deal with manifolds embedded in Euclidean space. The proofs of these theorems will be outlined. An algorithm for constructing these triangulations will also be described. Finally, directions for further work will be discussed.

2. SOME RIEMANNIAN GEOMETRY

A deep understand of Riemannian geometry will not be needed for the formulation of theorems in later sections. To learn more about Riemannian manifolds examine [3] or [6]. The essential features of a Riemannian manifold we will use are the abilities to measure angles and lengths of curves.

The curves we will be most interested in are geodesics. Essentially, geodesics are curves that locally minimize length. That is, they cannot be locally perturbed to reduce length. In particular, the shortest curve between two points is a geodesic. Distance in Riemannian manifolds is defined as the length of this shortest geodesic between two points. One difficulty is that there can be many geodesic between two points. See figure 1. In fact, there can be infinitely many shortest geodesics between two points. For example, examine antipodal points on a sphere. However, if all of our calculations are done in small enough neighborhoods we can avoid these problems. The concept of injectivity radius allows us to specify how close two points need to be to ensure that there is a unique shortest geodesic connecting them.

DEFINITION 2.1. *The injectivity radius at a point $x \in M$ is the largest radius, r , for which $B(x, r)$ is an embedded ball. The injectivity radius of M is the infimum of the injectivity radii at each point.*

An equivalent definition of the injectivity radius at $x \in M$ is half the length of the shortest closed geodesic going through




Figure 1: Three of the infinitely many geodesics between two points in a torus.

x. If the manifold has any interesting topology then it happens on a scale larger than the injectivity radius. To ensure there cannot be any non-trivial topology within a simplex of the triangulation, one of our density requirements will be that within each ball of some fixed radius smaller than the injectivity radius there is a point in the set.

To get an accurate representation of the manifold, we need to consider more than just the topology. We also need to represent the geometry accurately. In order to do this we need to be sure that areas of higher positive curvature have a higher density of points. The difficulty with large areas of positive curvature is that it becomes possible to have many geodesics connecting two points for geometric instead of topological reasons. To ensure that this does not cause a problem we need to work inside strongly convex balls, for inside them there is a unique geodesic connecting any two points.

DEFINITION 2.2. $C \subset M$ is said to be convex if given any two points $x, y \in C$ there is a unique shortest geodesic in C connecting x and y . $C \subset M$ is strongly convex if given any two points in C the unique shortest geodesic connecting them has its interior contained in C furthermore this is the only geodesic in C connecting the points.

For convexity there must be a shortest geodesic in the set connecting two points, however, there could be a shorter geodesic connecting them that leaves the set. For strong convexity to hold, however, this shortest geodesic in the manifold cannot leave the set. In Euclidean space every convex set is also strongly convex. For an example of a small ball that is convex but not strongly convex see figure 2. In this rounded off cone the disk is convex because there is a unique geodesic inside the disk connecting any two points. However, for some pairs of points there is a shorter geodesic connecting them that goes the other way around the cone, so the disk is not strongly convex.

To avoid problems with balls that are not embedded we work inside the injectivity radius. There are similar radii to ensure that a small ball will be convex or strongly convex. The main reason we want to work in strongly convex neighborhoods is that we know no shortest geodesic between two points ever leaves it. This allows us to ignore what happens in the manifold outside of the ball.

DEFINITION 2.3. Given $x \in M$. Then the convexity ra-




Figure 2: A convex set that is not strongly convex.

dius (respectively strong convexity radius) of M at x is the largest $r > 0$ for which $\overline{B(x, r)}$ is convex (strongly convex, respectively). The convexity or strong convexity radius of M is the infimum of the radii at all points in M .

Working in small convex balls will eliminate many of our problems. However, we need to know that they exist and their radii are strictly positive. In any compact Riemannian manifold we can construct a lower bound on these radii. This lemma is proved by comparing M to a sphere with strictly larger curvature. Many arguments in Riemannian geometry use a similar technique of comparing the manifold to one with constant curvature. Upper bounds on length are often obtained by comparing to a sphere with larger curvature and lower bounds are obtained by examining hyperbolic space with smaller curvature.

LEMMA 2.4 ([3]). If κ is a positive upper bound on the sectional curvature of M then

1. the convexity radius of M is at least $\min\{\text{inj}(M), \frac{\pi}{2\sqrt{\kappa}}\}$
2. similarly, the strong convexity radius of M is at least $\min\{\frac{\text{inj}(M)}{2}, \frac{\pi}{2\sqrt{\kappa}}\}$.

These ideas of injectivity and convexity radii are the key concepts from Riemannian geometry that we will use to construct our triangulations. With them we can do all of our calculations within small balls. We will see that inside these balls the geometry does not deviate enough from Euclidean to cause difficulties.

3. DELAUNAY TRIANGULATIONS

The definition of Delaunay triangulations is the same in Riemannian geometry as it is in \mathbf{R}^n . They are defined as having the empty circumscribing sphere property: the minimal radius circumscribing sphere for any simplex contains no vertices of the triangulation in its interior. However, there are several possible problems with this definition. How do we know the circumscribing sphere is unique and embedded? We need to deal with these issues to ensure that the property is well defined. This is where we rely on having a good sampling of points to use to construct the triangulation.

The existence of Delaunay triangulations in two dimensions was proven by the first author in [8]. However, the techniques used relied heavily on the dimension. Different techniques need to be used in higher dimensions. First we determine on how small of a scale we need to work to ensure that we can minimize complications due to the non-Euclidean geometry. The first requirement will be that all simplices of our Delaunay triangulation and their neighbors are contained inside strongly convex balls. Otherwise, we would run into problems with the triangulation being well defined.

DEFINITION 3.1. For any $x \in M$ define the density radius, $\text{rad}(x)$, to be $\frac{1}{5}$ the strong convexity radius of M at x .

The density radius tells us what scale we need to work on around any given point. However, problems can arise if the scale grows rapidly as locations change. So our density requirements include a condition that the density radius changes in a controlled manner.

DEFINITION 3.2. A finite set of points $X \subset M$ is said to be sufficiently dense if for every $y \in M$ and $z \in B(y, 4 \text{rad}(y))$ the ball of radius $\text{rad}(y)$ centered at z contains a point of X in its interior.

This definition enables us to ensure that the sampling of points is dense everywhere and that there are no major changes in the geometry as we move small distances in the manifold. While this density requirement is somewhat cumbersome, it allows local control over the number of points needed to be sufficiently dense. If global control rather than local is desired there is a much simpler way to guarantee a fine enough distribution of points for the Delaunay triangulation to exist.

LEMMA 3.3. If κ is a positive upper bound on the sectional curvature of M and $r = \min\{\frac{\text{inj}(M)}{10}, \frac{\pi}{10\sqrt{\kappa}}\}$ then if every ball of radius r in M contains a point in X then the points are sufficiently dense.

This stronger density condition will often be easier to use due to its simplicity. However, the local conditions will often be desirable as they require far fewer vertices in the triangulation. With either of these density conditions, we can guarantee that the Delaunay triangulation exists.

A second condition is that the set of points is generic: if M is d dimensional then $d + 2$ of the points never lie on the boundary of a round ball. Without this condition, instead of obtaining a triangulation, we obtain a cell complex.

THEOREM 3.4 ([9]). If $X = \{x_1, \dots, x_n\} \subset M$ is a generic, sufficiently dense set of points then there exists a unique Delaunay triangulation with vertices x_1, \dots, x_n .

COROLLARY 3.5 ([9]). For a sufficiently dense set of points $X \subset M^d$, any $d+1$ points, $x_1, \dots, x_{d+1} \subset B(z, \text{rad}(z)) \cap$




Figure 3: A Voronoi diagram when the density of points is too low.

X for any $z \in M$, determine a unique minimal radius circumscribing sphere.

Essentially this corollary says that $d + 1$ close enough points determine a unique d -sphere of small radius. This corollary ensures that the notion of being a Delaunay triangulation is well-defined.

The definition of the triangulation only deals with the combinatorics of the triangulation. From a topological point of view this is enough, but in some contexts we will want to know how the simplices and their faces are structured geometrically. In Euclidean space the faces of the simplices are planes. In Riemannian geometry the concept of a plane is not as simple to use. For our purposes we will define a plane in a Riemannian manifold as follows. Given two points $x, y \in M$ and a constant c then a plane will be $\{z \in M : [d(x, z)]^2 - [d(y, z)]^2 = c\}$. These planes are known as middle planes. Note that in Euclidean space this is equivalent to the usual definition of a plane. Inside small balls these planes have many of the same properties as planes in Euclidean space. The faces of simplices of the triangulation will consist of these planes. See [9] for more details.

4. VORONOI DIAGRAMS

Voronoi diagrams in Riemannian manifolds are defined exactly the same as in Euclidean space. They also have most of the same properties.

DEFINITION 4.1. If $X = \{x_1, \dots, x_n\} \subset M$ then the Voronoi diagram for X divides M into regions V_i for each x_i defined as follows:

$$V_i = \{z \in M : d(x_i, z) \leq d(x_j, z) \forall j\}$$

If Euclidean space the Voronoi diagram is always a cell complex. However, this will not be the case in general, see figure 4 to see how the Voronoi regions in a torus might not be disks. However, if the same density requirements are met as in the previous section then it will form a cell complex dual to the Delaunay triangulation. Observe that the faces of each of the regions of the Voronoi diagram are planes as defined in the previous section.




Figure 4: The medial axis for a manifold in \mathbf{R}^2 .

THEOREM 4.2 ([9]). *For a sufficiently dense set of points then Voronoi diagram is a cell complex for M . Furthermore, the Voronoi diagram is dual to the Delaunay triangulation.*

5. SUBMANIFOLDS OF EUCLIDEAN SPACE

Any manifold embedded in Euclidean space inherits a Riemannian structure where the lengths of curves on the manifold are the Euclidean lengths of the curves. When working with submanifolds of \mathbf{R}^n the proofs to many of the theorems in the previous section become much easier. However, this does not put any restrictions on the manifolds that can be considered. In 1956, Nash [11] proved that any Riemannian manifold can be embedded in Euclidean space so that its metric is induced by the embedding.

There is one further requirement for a sample of points when working in Euclidean space. Depending on the embedding it is possible for two points to be much closer in Euclidean space than they are in the manifold. This can be dealt with by knowing which points in Euclidean space have a unique closest point in the manifold.

DEFINITION 5.1. *The medial axis for a manifold M embedded in \mathbf{R}^n is the closure of the set of points with more than nearest neighbor in M .*

The additional requirements to ensure that the triangles can be replaced by flat ones are similar to those used in the surface reconstruction techniques used in [1].

DEFINITION 5.2. *Given a submanifold M of \mathbf{R}^n the local feature size at $x \in M$, denoted $\text{LFS}(x)$, is the distance from x to the medial axis for M .*

THEOREM 5.3. *If $M \subset \mathbf{R}^n$ is a submanifold of any dimension and X is a finite set of points so that*

1. X is sufficiently dense
2. for every $y \in M$, the ball of radius $\frac{1}{4}\text{LFS}(y)$ contains a point of X in its interior

then the Euclidean approximation of M from the Delaunay triangulation is ambiently isotopic to M . Furthermore, this isotopy induces a homeomorphism between M and its approximation.

If M' is the piecewise linear approximation of M then the homeomorphism takes $x \in M'$ to the closest point to x in M . This map is well defined since the above assumptions ensure that the approximation is disjoint from the medial axis. The ambient isotopy between the two surfaces follow the segments connecting points in M' to their images under the homeomorphism.

6. SKETCH OF PROOFS

The existence of Delaunay triangulations will be established by first showing that for a sufficiently dense set of points the Voronoi diagram is a cell decomposition of the manifold. For a generic set of points we will show that the Voronoi diagram is dual to a triangulation, in fact the Delaunay triangulation.

Let $X = \{x_1, \dots, x_n\}$ be a sufficiently dense set of points in M . Define $f : M \rightarrow \mathbf{R}^n$ by

$$f(p) = (d(x_1, p), \dots, d(x_n, p))$$

We will use f to study the Voronoi diagram.

LEMMA 6.1. *f is an embedding.*

PROOF. First observe that no matter which direction you move in M , the density conditions ensure that you must get further from some x_i . Pick any $p \in M$ for which $f(p) = f(q)$ for some $q \in M$ and let $d_i = d(p, x_i)$. Define a sets $I(a)$ as follows:

$$I(a) = \bigcap_{i=1}^n B(x_i, d_i + a)$$

Notice that $p \in I(0)$. Also, since small balls are strongly convex it can be shown that $I(a)$ is also strongly convex for sufficiently small a . If f is not an embedding, then $I(0)$ contains more than one point. By strong convexity, a geodesic arc connecting any two points must have its interior contained in the interior of $I(0)$. This implies for some $a < 0$ that $I(a) \neq \emptyset$. So it is possible to move from p along a geodesic and enter $I(a)$. If this could happen, then the distance to every x_i would decrease. This is impossible since the distance to one of the points must increase, so f must have been an embedding. \square

Note that $f(M)$ is contained in the positive orthant and that there is a natural cell decomposition of it into infinite simplicies

$$C_i = \{(z_1, \dots, z_n) : 0 \leq z_i \leq z_j \forall j\}$$

Notice that the Voronoi region V_i are precisely $f^{-1}(C_i)$. We will study these Voronoi regions by examining how $f(M)$ meets these simplicies in \mathbf{R}^n .

LEMMA 6.2. *V_i is a ball.*

PROOF. Notice that $V_i \subset B(x_i, \frac{1}{5}\text{conv}(x_i))$, where $\text{conv}(p)$ is the strong convexity radius of M at p . If this were not the case then there would exist a point $p \in V_i$ with $d(x_i, p) > \frac{1}{5}\text{conv}(x_i)$ which would contradict the density conditions on the points.

Let $p \in V_i$. Since p is inside the strong convexity radius there exists a unique geodesic γ connecting p and x_i . Pick any point, q , in the interior of γ . We claim that q is in the interior of V_i . If not, $d(x_i, q) = d(x_j, q)$ for some $j \neq i$. The triangle inequality implies that $d(x_j, p) \leq d(x_j, q) + d(q, p) = d(x_i, q) + d(q, p) = d(x_i, p)$. Hence $d(x_j, p) = d(x_i, p)$ and the unique geodesics connecting p to x_i and from p to x_j coincide on an interval. This only happens if all of the points are on the same geodesic, hence $x_i = x_j$. This shows that V_i is star convex from x_i . Since each ray from x_i hits the boundary in at most one point and that x_i is in the interior of V_i , V_i must be a ball. \square

To show that the Voronoi diagram is a cell complex, we need to examine how these balls intersect. To understand this intersection we need to restrict how $M' = f(M)$ intersects faces of the cell decomposition of \mathbf{R}^n and parallel copies of the faces. Let z_1, \dots, z_n be the coordinates in \mathbf{R}^n .

LEMMA 6.3 (TRANSVERSALITY). *If $p \in M$ and we can reindex so that $x_1, \dots, x_k \in X$ are the points of X inside $B(p, 3 \text{ rad}(p))$ then any non-zero vector in the span of z_1, \dots, z_k is not in $T_{f(p)}M'$, the tangent plane to M' at $f(p)$.*

PROOF. Notice that if the lemma were false then it would be possible to travel in some direction from p and the distance from x_1, \dots, x_k would change, but the distances to any other x_i would remain fixed. But by our density assumptions, it can be shown that movement in any direction will cause the distance to some $x_i, i > k$ to decrease. \square

To show that the Voronoi diagram is a cell complex we will show that each cell is of a particularly nice type called normal. These definitions are generalizations of the theory of normal surfaces in the 3-manifold topology.

DEFINITION 6.4. *A normal d -ball B in a convex Euclidean n -cell C is a ball such that B intersects a k dimensional face of C in a single unknotted $(k - n + d)$ -ball or the intersection is empty. See figure 6.*

Normal balls have been studied in the generalizations of normal surface theory and various existence results proven, though not published. The following theorem allows us to prove that each of the Voronoi cells is a normal ball. A proof of the theorem can be found in [9].

THEOREM 6.5. *If C is a convex Euclidean n -cell and $B \in C$ is a d -manifold such that*

1. B intersects each k dimensional face of C in an $(k - n + d)$ -manifold




Figure 5: Normal arcs and disks in a tetrahedron.

2. If F is a face of C then $F \cap B$ intersects every hyper-plane in F transversely.

then B is a normal ball.

LEMMA 6.6. V_i is a normal d -ball in the cell C_i .

PROOF. Since the set of points was chosen to be generic the first requirement of the theorem above is satisfied. And the transversality lemma ensures that every face M' meets it in transversely. So the theorem can be applied to show that M' hits each of the C_i in a union of normal balls. But we know V_i is connected, so it is a single normal ball. \square

The intersection of any two Voronoi cells occurs in a face of the cell decomposition of \mathbf{R}^n . The above lemma says that these intersections are always a single ball or empty. This is the requirement for the Voronoi diagram to be a cell complex.

If we perturb our points slightly we can ensure that the Voronoi diagram is dual to a triangulation. The same arguments as used in Euclidean space can be applied to show that the triangulation dual to the Voronoi diagram is in fact the Delaunay triangulation. This method is useful in proving that Delaunay triangulations exist, however working in such a large dimensional Euclidean space is completely impractical. In the next section we will discuss other means for construction this triangulation now that we know it exists.

7. ALGORITHMS



We will discuss algorithms for constructing the Delaunay triangulation for a sufficiently dense set of points. We will assume the existence of two oracles that allow us to do the necessary calculations on the Riemannian manifold. The first will be for measuring the distance between any two points on the manifold. The second will find the center of the smallest radius sphere through a given set of points. Using these two oracles, the construction algorithm follows using an incremental algorithm to calculate the triangulation in the neighborhood of every point.

The incremental algorithm proceeds by adding a point at a time to the triangulation. Initially the triangulation near the first point is arbitrary. But a sequence of Pachner moves or flips can be performed upon any simplicies that do not

have empty circumscribing spheres. See figure 7 to see the moves in two and three dimensions. In Euclidean space it can be shown that this will always produce the Delaunay triangulation.

ALGORITHM 7.1. Let x_1, \dots, x_n be a sufficiently dense set of points in M .

1. For each i construct the Delaunay triangulation T_i for all simplices meeting x_i .
 - (a) Start with a d -simplex made from x_i and its d nearest neighbors.
 - (b) Add each of the x_j incrementally.
 - (c) If x_j is in the interior of an existing simplex subdivide that simplex by adding the new vertex. Otherwise, add simplices by coning the faces visible from x_j to the new vertex x_j . See figure 7a.
 - (d) If any of the simplices adjacent to x_i do not have empty circumscribing spheres, perform flips until they do so. See figure 7b.
 - (e) Repeat (b) - (d) until all points have been added. Let T_i be the triangulation of the star of x_i .
2. Piece together the triangulations, T_i , to obtain the Delaunay triangulation.



While the incremental algorithm cannot be used to produce a Delaunay triangulation for the entire manifold, it does construct the correct triangulation for the simplices adjacent to the initial point. In step 2, it can be shown that the T_i agree on their overlap. The proof of these two statements can be seen in [10].

The second algorithm we want to discuss involves constructing triangulations that accurately approximate the Riemannian metric. Our Delaunay triangulation provides an approximation of the manifolds geometry as follows. Recall the map $f : M \rightarrow \mathbf{R}^n$ in section 6. A piecewise flat manifold N can be obtained by using Euclidean simplices in \mathbf{R}^n to construct the triangulation. There is a natural homeomorphism $g : N \rightarrow M$. We will define a measure to evaluate how closely the Delaunay triangulation represents the geometry. The measure uses a bi-Lipschitz relation comparing the measurements of lengths of curves and angles.

In Riemannian manifolds, length and angles are measured using an inner product on the tangent space at any point. Using the piecewise flat structure on our approximation, there is an inner product that allows us to also make these calculations on the approximation. The precise definition is technical and will be omitted, but we will say that the approximation is ϵ -close to the Riemannian manifold if evaluation of the two inner products never differ by more than a factor of ϵ . However, this notion of ϵ -close has a nice interpretation in the measurements of lengths and angles. And for our purpose, we can use this as the definition of ϵ -close.

LEMMA 7.2. If N is ϵ -close to $M = g(N)$ then




Figure 6: The Pachner moves in two and three dimensions.




Figure 7: (a) Adding a new vertex to the triangulation and (b) do flips so that the triangulation is Delaunay.

1. if γ is any curve in N and $\gamma' = g(\gamma)$ then

$$(1 - \epsilon)\text{length}(\gamma) < \text{length}(\gamma') < (1 + \epsilon)\text{length}(\gamma)$$
2. if θ is the angle of intersection of two curves in N and θ' is the angle of intersection of the same curves in M then $|\cos \theta - \cos \theta'| < \epsilon$.

The density requirements for the existence of Delaunay triangulations do not take into account regions of negative curvature. This does not affect getting a topologically accurate triangulation; however, they do affect how well the geometry is represented. To deal with this we will need a stronger sampling requirement.

DEFINITION 7.3. A sufficiently dense set of points, $X = \{x_1, \dots, x_n\}$, is δ -sampled if for every point $p \in M$, we have that $B(p, \delta \frac{\pi}{10\sqrt{\kappa_p}})$ contains one of the x_i , where κ_p is a bound on the absolute value of the sectional curvature at p .

THEOREM 7.4. For every $\epsilon > 0$ there exists a value for δ such that the Delaunay triangulation of any manifold using any δ -sampled set of points will be ϵ -close to the Riemannian metric. Furthermore, δ only depends on ϵ and the dimension of the manifold.

This theorem ensures that if you subdivide your triangulation enough you get can arbitrarily accurate representations of the geometry of the manifold. However, the required density of points might not be provided. For point samples that are 1-dense or better, we can measure the accuracy of the approximation. If we do not have the desired accuracy then

we can add Steiner points until it is reached. This gives an algorithm to constructing arbitrarily precise triangulations.

ALGORITHM 7.5. Start with a 1-dense set of points in M .

1. Construct the Delaunay triangulation.
2. For each simplex of the triangulation, bound the sectional curvature by examining the tangent planes at each of the vertices.
3. By comparing the curvature bounds to the lengths of the edges, it can be seen if the sample of points is dense enough locally.
4. If there is too much deviation for the estimate to be ϵ -close inside the tetrahedra then add a Steiner point at the center of the circumscribing sphere and retriangulate.
5. Repeat steps 2-4 until the desired accuracy is reached.

This algorithm will terminate with a triangulation of M that represents the topology and geometry of the manifold accurately. The number of points added will be large if the manifold has large regions of negative curvature.

8. AN EXAMPLE

As an example for the results of these algorithms, we will examine a piece of the surface $z = x^2 + y^2$. See figure 8a to see the density requirements to guarantee a certain quality of approximation. The ellipses in the plane are projections of round circles in the surface; each must have a vertex of the triangulation in its interior to meet the density requirements. Where the surface has higher curvature a higher density of points is required, and in the direction of greatest change the density is highest. Figure 8b shows the resulting triangulation. The triangles vary in size and the vertices are positioned to ensure a high quality approximation. To achieve a similar quality triangulation using triangles of uniform size requires approximately two and a half times the number of vertices.

9. CONCLUSIONS AND FUTURE WORK

In this paper we have introduced the concept of Delaunay triangulations to Riemannian geometry and found conditions to ensure their existence. The algorithms discussed allow arbitrarily accurate approximations of a manifold using a variety of quality measures. At this point these algorithms are still theoretical; however, with improved heuristics they could become practical.

There are many possible applications to these theorems. One of them is surface reconstruction using tangent plane information. Note that the theorems work for any dimension manifold embedded in any dimension Euclidean space. This opens up the possibilities for algorithms to reconstruct curves in \mathbb{R}^3 and possible applications to manifold learning where the dimensions of spaces can be quite large.

Another application is in anisotropic meshing for modelling solutions to complex differential equations. The solutions

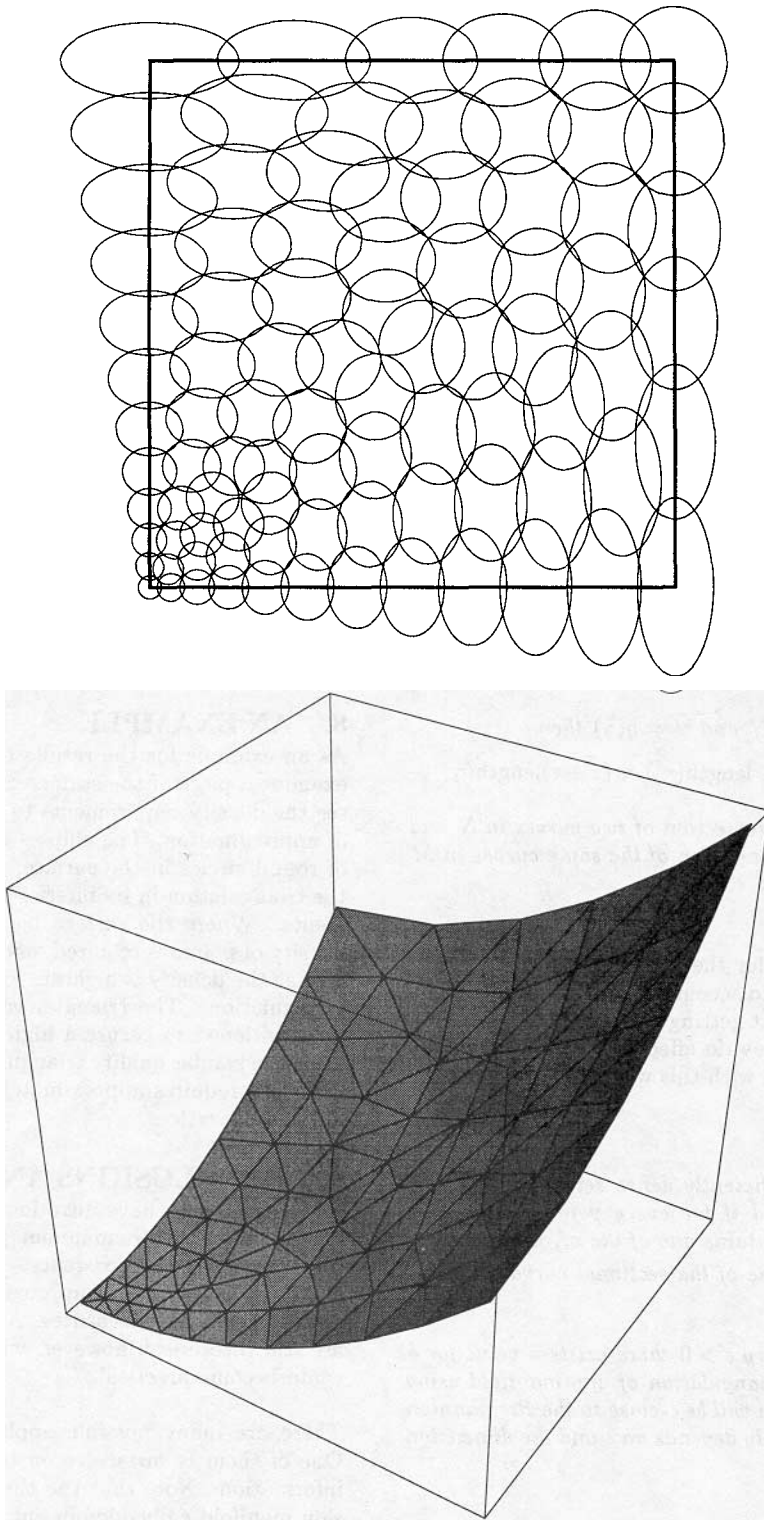


Figure 8: Example of a Delaunay triangulation for a Riemannian manifold in \mathbb{R}^3

to these equations can be thought of a submanifold of some Euclidean space and the algorithms discussed here can be used to approximate these manifolds accurately.

A final application is to understand triangulations of 3-dimensional manifolds. In 3-dimensional topology the study of geometric structures on manifolds is an important area of research. The techniques in this paper can be used to construct triangulations for these manifolds that are compatible with their geometric structures. These triangulations have nice combinatorial properties and in many cases are conjectured to be minimal. This could lead to some interesting results in topology.

10. ACKNOWLEDGEMENTS

We would like to thank Herbert Edelsbrunner for helpful discussions about the algorithms, Hyam Rubinstein for his suggestions for dealing with normal submanifolds of Euclidean space and Ben Burton for his comments.

11. REFERENCES

- [1] N. Amenta and M. Bern. Surface Reconstruction by Voronoi Filtering. *Discrete Comput. Geom.* **22** (1999) 471-504.
- [2] M. Bern and D. Eppstein. Mesh Generation and Optimal Triangulation. *Computing in Euclidean Geometry*, edited by F. K. Hwang and D.-Z. Du, World Scientific, 1992, 47-123.
- [3] I. Chavel. *Riemannian Geometry: A Modern Introduction*. Cambridge University Press, 1993.
- [4] P. Chew. Guaranteed-Quality Mesh Generation for Curved Surfaces. *Proc. 9th Symp. Computational Geometry*, 1993, 274-280.
- [5] H. Edelsbrunner. *Algorithms in Combinatorial Geometry*. Springer-Verlag, Heidelberg, Germany, 1987.
- [6] W. Klingenberg. *Riemannian Geometry*. de Gruyter, Berlin, 1982.
- [7] C. Lawson. Generation of a Triangular Grid with Applications to Contour Plotting. Memo 299. Jet Propulsion Laboratory, Pasadena, CA, 1972.
- [8] G. Leibon. *Random Delaunay Triangulations, the Thruston-Andreev Theorem, and Metric Uniformization*. Ph.D. Thesis, UCSD, 1999.
- [9] G. Leibon and D. Letscher. Delaunay Triangulations for Riemannian Manifolds. Part I: Existence. Submitted for publication, 2000.
- [10] D. Letscher. Delaunay Triangulations for Riemannian Manifolds. Part II: Algorithms. In Preparation, 2000.
- [11] J. Nash. The Imbedding Problem for Riemannian Manifolds. *Ann. of Math.* **63** (1956), 20-63.