Uniform Samples of Generic Surfaces Have Nice Delaunay Triangulations*

Jeff Erickson†

University of Illinois at Urbana-Champaign
jeffe@cs.uiuc.edu
http://www.cs.uiuc.edu/~jeffe

Submitted to SOCG 2003 — December 4, 2002

Abstract

Let Σ be a fixed smooth surface in $\mathbb{R}^3$, such that no medial ball touches Σ more than four times, counting with multiplicity, or more than three times at any single point. We show that the Delaunay triangulation of any uniform sample of Σ has complexity $O(n \log n)$ in the worst case. We also prove that the Delaunay triangulation of n random points on Σ has complexity $O(n \log^3 n)$ with high probability. In both upper bounds, the hidden constants depend on the surface.

I'm checkin' 'em out; I'm checkin' 'em out
I've got it figured out; I've got it figured out
There's some good points, some bad points,
But it all works out - I'm just a little freaked out.

— Phish, "Cities", *Slip Stitch & Pass* (1997)
after Talking Heads, "Cities", *Fear of Music* (1979)

Since we cannot hope for order, let us withdraw with style from the chaos.

— Lord Malquist, *Lord Malquist and Mr. Moon*, by Tom Stoppard (1966)

*See http://www.cs.uiuc.edu/~jeffe/pubs/smooth.html for the most recent version of this paper.

†Partially supported by a Sloan Foundation Fellowship, NSF CAREER award CCR-0093348, and NSF ITR grants DMR-0121695 and CCR-0219594.

1 Introduction

Delaunay triangulations and their dual Voronoi diagrams are among the most commonly used and thoroughly studied structures in combinatorial geometry. One application that has received considerable attention recently is curve and surface reconstruction [1, 2, 3, 4, 12, 18, 19, 23, 29, 30]. The input to the surface reconstruction problem is a set of unorganized points from an unknown surface Σ in $\mathbb{R}^3$, and the goal is to construct a geometric approximation of Σ with the correct topology. Most recent reconstruction algorithms begin by constructing the Delaunay triangulation or Voronoi diagram of the input points. (The recent work of Dey, Funke, and Ramos [18, 23] is a notable exception.) Since three-dimensional Delaunay triangulations can have complexity $\Omega(n^2)$ in the worst case, these algorithms have worst-case running time $\Omega(n^2)$. However, this behavior is almost never observed in practice [12, 17] except for highly-contrived inputs [21]. For all practical purposes, Delaunay triangulations of surface points appear to have linear complexity.

The first subquadratic complexity bound for Delaunay triangulations of surface points was obtained by Golin and Na [25, 26, 27]. They proved that if n points are chosen uniformly at random on the surface of any fixed convex polytope in $\mathbb{R}^3$, the expected complexity of their Delaunay triangulation is $O(n)$. Using similar techniques, they recently showed that a random sample of a fixed *nonconvex* polyhedron has Delaunay complexity $O(n \log^4 n)$ with high probability [28]. In fact, their analysis applies to any fixed set of triangles in $\mathbb{R}^3$.

Attali and Boissonnat [9] recently proved that the Delaunay triangulation of any (ε, k)*-sample* of a fixed polyhedral surface has complexity $O(k^2 n)$, improving their previous upper bound of $O(n^{7/4})$ (for constant k) [8]. A set of points is called an (ε, k)-sample of a surface Σ if every ball of radius ε whose center lies on Σ contains at least one and at most k points in P. A simple application of Chernoff bounds implies that a random sample of n points on a fixed surface is an $O(\varepsilon, O(\log n))$-sample with high probability, where $\varepsilon = O(\sqrt{(\log n)/n})$. (See Theorem 4.2 for a similar derivation.) Thus, Attali and Boissonnat's result improves Golin and Na's high-probability bound for random points to $O(n \log^2 n)$.

The hidden constants in all these bounds depend on geometric parameters of the fixed surface, such as the number of facets, angles between adjacent edges, and angles between facet planes. For this reason, none of these bounds apply to smooth surfaces.

Previously known bounds for non-polyhedral surfaces are much weaker. In two earlier papers [21, 22], we analyzed the complexity of three-dimensional Delaunay triangulations in terms of a geometric parameter called the *spread*, defined as the ratio between the largest and smallest pairwise distances. Our results imply that any (ε, k)-sample of any fixed (not necessarily polyhedral or smooth) surface Σ has Delaunay complexity $O(k^2 n^{3/2})$. Moreover, this bound is tight in the worst case, at least for $k = O(1)$; a right circular cylinder with constant height and radius has a uniform $(\varepsilon, O(1))$-sample with Delaunay complexity $\Omega(n^{3/2})$. However, this surface is extremely degenerate; every medial ball intersects the surface in an infinite number of points.

The (ε, k)-sampling condition limits *oversampling*; without some limit of this form, Delaunay triangulations can be arbitrarily complex. Specifically, for any surface other than the sphere or the plane and any sampling density $\varepsilon > 0$, there is an ε-sample whose Delaunay triangulation has complexity $\Theta(n^2)$, where n is the minimum number of sample points in any ε-sample [21].

In this paper, we show that under a mild uniform sampling condition, the Delaunay triangulation of a set of points on a fixed *generic smooth* surface in $\mathbb{R}^3$ without boundary is only $O(n \log n)$ in the worst case. Loosely, a smooth surface is generic if no medial ball touches it more than four times, counting with multiplicity, or more than three times at a single point. We also show that the Delaunay triangulation of n random points on a fixed generic smooth surface has complexity

$O(n \log^3 n)$ with high probability. Like all previous subquadratic upper bounds for surface samples, the hidden constants in both our bounds depend on geometric parameters of the fixed surface Σ.

2 Background

In this section, we provide necessary background information and develop some useful tools for the proof of our main result.

2.1 Medial Balls and Medial Axes

Every point on the surface is a point of tangency for two medial balls, one on each side of the surface. For points on the convex hull of the surface, one of the medial balls degenerates to a closed halfspace whose boundary plane is tangent to the surface at (generically) at most three points. To avoid these boundary cases, we will directly consider only *interior* medial balls, which are centered inside the surface Σ.[1]

Like most previous results about Delaunay triangulations of surface points, our proof boils down to the following intuitive observation: As the density of samples increases, Delaunay balls intersect less and less of the surface. In the limit, Delaunay balls approach balls that intersect Σ only at their boundary, generically in at most four points. To exploit this observation, we need to understand exactly how Delaunay balls approach their limiting behavior.

2.2 Smooth Functions and Surfaces

In this paper, a *smooth* function or surface is C^4 and real-analytic; that is, partial fourth derivatives exist[2] and Taylor series approximations have a positive radius of convergence at every point on the surface. As pointed out by Choi *et al.* [16], there are C^∞ curves and surfaces whose medial axes are infinitely complex. For example, the medial axis of the curve $y = e^{-1/x^2} \sin^2(1/x)$ has an infinite number of branches; moreover, the local feature size is positive at every point on the curve.

Our analysis makes heavy use of the following standard result about Taylor-Maclaurin series [33]:

Taylor's Theorem. *For any integer $k > 0$ and any C^k function $f : \mathbb{R} \to \mathbb{R}$, we have*

$$f(x) = \sum_{i=0}^{k-1} f^{(i)}(0)\, \frac{x^i}{i!} + f^{(k)}(x^*)\, \frac{x^k}{k!}$$

for some $0 < x^* < x$.

If f is smooth and x is sufficiently small (within the radius of convergence of the Taylor series), the last term in this sum converges to 0 as $k \to \infty$ or as $x \to 0$. Whenever we apply Taylor's theorem, we assume that k is sufficiently large and x sufficiently small that the absolute value of the error term $f^{(k)}(x^*)\, x^k/k!$ is at most $|f(x)|/2$. In most cases, we will use the smallest value of k such that $f^{(k-1)}(0) \neq 0$.

---

[1] Alternatively, we could avoid these boundary conditions by applying the standard stereographic lifting transformation, which maps Delaunay (and anti-Delaunay) edges to convex hull edges and medial (and anti-medial) balls to supporting halfspaces. Our main result can be interpreted in this context as follows: The convex hull of a uniform sample of a fixed smooth 2-manifold subset of the unit sphere in $\mathbb{R}^4$ has complexity $O(n \log n)$.

[2] Actually, we only need fourth derivatives at extreme points of principal curvature, and then only in the principal curvature direction; otherwise, we only assume the surface is C^2.

Taylor's theorem can be generalized to multivariate functions by expanding one variable at a time. For example, for any smooth function $f : \mathbb{R}^2 \to \mathbb{R}$, and any integer $k > 0$, Taylor's theorem implies that

$$f(x, y) = \frac{\partial^i f}{\partial x^i}(0, y)\, \frac{x^i}{i!} + \frac{\partial^k f}{\partial x^k}(x^*, y)\, \frac{x^k}{k!}$$

for some $0 < x^* < x$. Applying Taylor's theorem with respect to y to each term in this summation gives us, for any integers $l_0, l_1, \ldots, l_{k-1} > 0$, the identity

$$f(x, y) = \sum_{i=0}^{k-1} \left(\sum_{j=0}^{l_i - 1} \frac{\partial^{i+j} f}{\partial x^i\, \partial y^j}(0, 0)\, \frac{x^i\, y^j}{i!\, j!} + \frac{\partial^{i+l_i} f}{\partial x^i\, \partial y^{l_i}}(0, y_i^*)\, \frac{x^i\, y^{l_i}}{i!\, l_i!} \right) + \frac{\partial^k f}{\partial x^k}(x^*, 0)\, \frac{x^k}{k!} + \frac{\partial^{k+1} f}{\partial x^k\, \partial y}(x^*, y_k^*)\, \frac{x^k\, y}{k!}$$

for some $0 < y_0^*, y_1^*, \ldots, y_k^* < y$. We will use this version of Taylor's theorem, usually without comment, to approximate functions by their smallest nonzero partial derivatives, by omitting the error terms (those involving x_* or some y_i^*). Whenever we apply this approximation, we will assume the arguments x and y are sufficiently small that the cumulative approximation error is less than a constant factor.

2.3 Generic Medial Balls

Our results rely on a classification of medial balls for generic smooth surfaces, and their surface contacts, first announced Bryzgalova in 1977 (in a more general form) [15]. See also Mazov [31, 32]; Bruce, Giblin, and Gibson [14]; Bogaevsky [10, 11]; Anoshkina *et al.* [5]; and Giblin and Kimia [24]. For further background on singularity theory, we refer the interested reader to the short survey by Arnol'd [6] or more detailed expositions by Bruce and Giblin [13] and Arnol'd *et al.* [7].

There are exactly two generic types of contact between a medial ball and a smooth surface in $\mathbb{R}^3$. An $\boldsymbol{A_1}$ contact occurs when the surface is tangent to the medial ball, but shares no higher-order derivatives. An $\boldsymbol{A_3}$ contact occurs where the surface is tangent to the ball and the radius of the ball equals the minimum principal curvature radius (with the appropriate sign). A_3 contact points are also *ridge* points, where the one of the principal curvatures is maximized. The A_3 contact points form a family of curves on the surface; topologically, each A_3 curve is either a circle or a closed interval. Every pair of A_3 curves is disjoint, even at their endpoints.

Similarly, there are exactly five types of medial balls—A_1^2, A_1^3, A_3, A_1^4, and A_1A_3—named according to the number and orders of their contact points. Each type of medial ball corresponds to a different type of feature of the medial axis; these are illustrated in Figure 1.

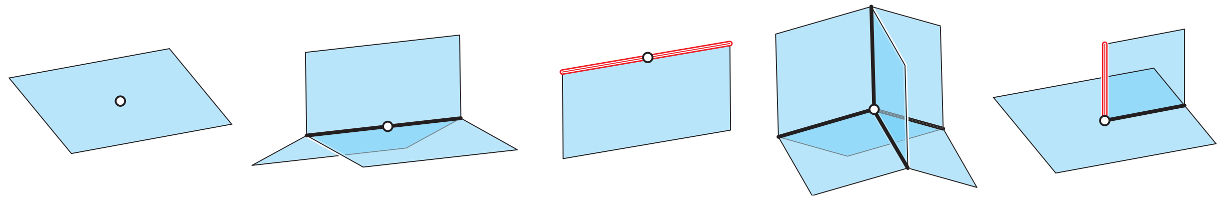

Figure 1. Generic medial axis features: A_1^2 (surface patch), A_1^3 (interior edge), A_3 (boundary edge), A_1^4 (interior vertex), and A_1A_3 (boundary vertex).

3 The Main Result

For the remainder of the paper, we fix a generic smooth surface Σ without boundary in $\mathbb{R}^3$. Because the Σ is fixed, **any finite, positive function of $\boldsymbol{\Sigma}$ is a constant**. Specific functions of Σ that appear in our analysis include the surface area, the minimum curvature radius, the minimum local feature size, the ratio between the largest enclosing ball and the smallest enclosed ball, the radius of convergence of Taylor series approximations, and most importantly, the 'genericity' of the surface. These functions of Σ affect both the hidden constant in the final $O(n \log n)$ upper bound and the minimum value of n to which the upper bound applies. Moreover, we have not even attempted to minimize the cumulative effects of these constants; even for very simple, well-behaved surfaces, the constant factors are likely to be enormous.

Following Attali and Boissonnat [9], we call a set P of points an (ε, k)-*sample* if any ball of radius ε centered on Σ contains at least one and at most k points in P. A simple packing argument implies that the number of points in any (ε, k)-sample of Σ is between $\Omega(1/\varepsilon^2)$ and $O(k/\varepsilon^2)$, where the hidden constants depend on the surface area and the maximum curvature of Σ.

Main Theorem. *The Delaunay triangulation of any (ε, k)-sample of any fixed generic smooth surface Σ has complexity $O(k^2 n \log n)$, where the hidden constant depends on Σ.*

Let $P = P(\varepsilon)$ be an (ε, k)-sample of Σ. To prove the Main Theorem, we will explicitly count only edges of the Delaunay triangulation of P. Euler's formula implies that any three-dimensional triangulation with n vertices and e edges contains at most $3e - 3n$ triangles and $2e - 2n$ tetrahedra, since the link of any vertex is a planar triangulation. Two points $p, q \in P$ are joined by an edge in the Delaunay triangulation of P if and only if they lie on the boundary of a closed ball B that excludes every other point in P. We call B a *Delaunay ball* for the edge pq.

We call a Delaunay edge *internal* if it has at least one Delaunay ball whose center is in the interior of Σ; otherwise, we call the edge *external*. Most of our analysis will be restricted to internal Delaunay edges. In Section 3.4, we count the external edges by exploiting the conformal invariance of Delaunay triangulations. Specifically, we apply a conformal transformation that turns the surface inside out, transforming the external Delaunay edges into internal Delaunay edges, except for $O(n)$ *boundary* edges, which we count with a simple packing argument.

We say that an internal edge pq is *local* if it has a Delaunay ball B such that p and q lie in the same component of $B \cap \Sigma$; otherwise, we call the edge *remote*. We count these two classes of edges separately in Sections 3.2 and 3.3. The number of local neighbors of a point depends on the shape of the intersection of Σ with a slightly enlarged medial ball. The number of remote neighbors depends on the differential behavior of the *medial reflection* function, which maps one contact point of an A_1^2 or A_1A_3 medial ball to the other contact point.

Finally, we split the points of Σ into three classes based on their distance to the A_3 contact curves, in terms of two small positive constants β_Σ and γ_Σ to be specified later. We say that a point is *good* if its distance to the nearest A_3 contact curve is greater than γ_Σ, *bad* if that distance is between $\beta_\Sigma\sqrt{\varepsilon}$ and γ_Σ, and *ugly* if that distance is less than $\beta_\Sigma\sqrt{\varepsilon}$. Intuitively, good points are 'obviously' A_1 contacts, ugly points are indistinguishable from A_3 contact points, and bad points interpolate between those two extremes.

3.1 Setup

The *ε-disk centered at p*, denoted $d(p, \varepsilon)$, is the intersection of Σ with a ball of radius ε centered at p. We call a ball *almost medial* if it is centered in the interior of Σ and it does not contain an ε-disk. Our sampling condition immediately implies that every ε-disk contains at least one point, which implies that any internal edge of the Delaunay triangulation of P has an almost medial Delaunay ball.

For any point $p \in \Sigma$, let $B_*(p)$ denote the interior medial ball that touches Σ at p. We respectively call the center $c_*(p)$ and radius $r_*(p)$ of $B(p)$ the *medial center* and *medial radius* of p.

Let $\kappa_1(p)$ and $\kappa_2(p)$ respectively denote the maximum and minimum principal curvatures of Σ at p, signed so that curvature toward the interior of the surface is positive. For every point p, we have $\kappa_2(p) \le \kappa_1(p) \le 1/r(p)$, and because Σ is generic, at least one of these two inequalities must be strict. The principal curvatures are equal only at isolated *umbilic* points, and $\kappa_1(p) = 1/r_*(p)$ if and only if p is an A_3 contact point.[3] We will write B_*, r_*, κ_1, and κ_2 when the point p is clear from context.

At each point p, we establish a local coordinate system where p is the origin, the tangent plane to Σ at p is the xy-plane, B_* is centered on the positive z-axis, the maximal curvature direction is the x-axis, and the minimum curvature direction is the y-axis. At umbilic points, we can choose the x- and y-axes arbitrarily.

3.2 Local Internal Edges

For any point p, let $L(p)$ denote the subset of Σ that contains all possible local internal Delaunay neighbors of p. We call $L(p)$ the *local neighborhood* of p. $L(p)$ clearly contains the disk $d(p, \varepsilon)$. Thus, a standard packing argument implies that the number of local internal neighbors of p is $O(k\,\mathrm{area}(L(p))/\varepsilon^2)$. In this section, we count the local internal edges by computing the area of the local neighborhood of each point in P.

Let B be an arbitrary almost-medial ball of radius r, let $p \in \Sigma$ be the point inside B and furthest from its boundary, and let δ denote the distance from from p to ∂B. Since B cannot contain the disk $d(p, \varepsilon)$, we must have $\delta < \varepsilon$. Moreover, B contains a ball B' that touches the surface only at p, so $r < r_*(p) + \delta < r_*(p) + \varepsilon$. Finally, let $h(p, \delta)$ denote the component of $B \cap \Sigma$ containing p.

In the local coordinate system at p, the sphere ∂B and the surface Σ can be approximately parameterized by the first few terms of their Taylor series:

$$\partial B(x, y) = -\delta + \frac{x^2}{2r} + \frac{y^2}{2r} + \cdots \qquad\qquad \Sigma(x, y) = \frac{\kappa_1 x^2}{2} + \frac{\kappa_2 y^2}{2} + \cdots$$

Thus, a first-order approximation of $h(p, \delta)$ is the set of points $\Sigma(x, y)$ where

$$x^2(1 - \kappa_1 r) + y^2(1 - \kappa_2 r) \le 2r\delta. \tag{1}$$

If $\kappa_2 \le \kappa_1 < 1/r$, then $h(p)$ is a closed elliptical disk whose major axis is horizontal and whose aspect ratio (width/height) is $\sqrt{(1 - \kappa_2 r)/(1 - \kappa_1 r)}$.

Lemma 3.1. *There are $O(kn)$ local internal edges with at least one good endpoint.*

Proof: Let p be a good point, let q be one of its local internal neighbors, and let B be a Delaunay ball for pq such that p and q lie on the boundary of the same component of $B \cap \Sigma$. If ε is sufficiently small, this component of $B \cap \Sigma$ is approximated by an ellipse with aspect ratio less than $\sqrt{(1 - \kappa_2 r_*)/(1 - \kappa_1 r_*)}$ that is tangent to p and q and does not contain an ε-disk. The local neighborhood $L(p)$ is the union of all such ellipses tangent to p. This union is an ellipse centered at p with height 4ε and width $4\varepsilon\sqrt{(1 - \kappa_2 r_*)/(1 - \kappa_1 r_*)}$. Thus, p has $O(k\sqrt{(1 - \kappa_2 r_*)/(1 - \kappa_1 r_*)})$ local internal neighbors. Let

$$\eta_\Sigma = \min_{\text{good points } p} \sqrt{\frac{1 - \kappa_2 r_*}{1 - \kappa_1 r_*}}$$

Because $\kappa_1 < 1/r_*$ except at A_3 contact points (which are not good), η_Σ is a finite positive constant. Every good point has $O(k\eta_\Sigma) = O(k)$ local internal neighbors, and there are trivially at most n good points. $\square$

[3] Contact between a medial sphere and an umbilic point is a so-called D_4 singularity, which can be transformed into an A_1A_3 singularity by perturbing the surface.

On the other hand, suppose $\kappa_1 \approx 1/r$, which can only happen if p is an A_3 contact point and $r \approx r_*$. Then the ellipse described by Equation (1) degenerates to two parallel lines. To accurately describe the shape of $h(p,\delta)$, we need to expand the Taylor series further. The maximum principal curvature κ_1 is locally maximized at p, so $\frac{\partial^3}{\partial x^3}\Sigma(0,0) = 0$, so we actually need the *fourth* partial derivative in the x-direction. We have

$$\partial B(x,y) \approx -\delta + \frac{x^2}{2r_*} + \frac{x^4}{8r_*^3} + \frac{y^2}{2r_*} + \cdots, \qquad \Sigma(x,y) = \frac{x^2}{2r_*} + \frac{\kappa_2 y^2}{2} + \frac{\lambda_1 x^4}{24} + \cdots,$$

where $\lambda_1 = \frac{\partial^4}{\partial x^4}\Sigma(0,0)$. Thus, a closer approximation of $h(p,\delta)$ is the set of points $\Sigma(x,y)$ such that

$$x^4\left(\frac{3-\lambda_1 r_*^3}{24 r_*^3}\right) + y^2\left(\frac{1-\kappa_2 r_*}{2r_*}\right) \le \delta \tag{2}$$

Because Σ is a generic surface, we must have $\lambda_1 r_*^3 < 3$, and by our earlier assumption, $\kappa_2 r < 1$. Thus, when p is an A_3 contact point, $h(p,\delta)$ is a centrally-symmetric convex oval that is wider along the x-axis than along the y-axis. The aspect ratio of this oval is

$$\left(\frac{24 r_*^3\delta}{3-\lambda_1 r_*^3}\right)^{1/4} \Big/ \left(\frac{2r_*\delta}{1-\kappa_2 r_*}\right)^{1/2} = \left(\frac{6r_*(1-\kappa_2 r_*)^2}{(3-\lambda_1 r_*^3)\delta}\right)^{1/4} \le \left(\frac{6r_*(1-\kappa_2 r_*)^2}{(3-\lambda_1 r_*^3)}\right)^{1/4}\frac{1}{\sqrt{\varepsilon}}$$

Lemma 3.2. *There are $O(k^2 n)$ local internal edges with at least one ugly endpoint.*

Proof: We define another constant

$$\zeta_\Sigma = \min_{A_3 \text{ contact points } p}\left(\frac{6r_*(1-\kappa_2 r_*)^2}{3-\lambda_1 r_*^3}\right)^{1/4}$$

Because our fixed surface Σ is generic, ζ_Σ is a positive and finite constant. The worst possible aspect ratio of the intersection oval of an almost-medial ball with Σ is approximately $\zeta_\Sigma/\sqrt{\varepsilon}$.

Let p be *any* point in P, let q be one of its local internal neighbors, and let B be a Delaunay ball for pq such that p and q lie on the boundary of the same component of $B\cap\Sigma$. If ε is sufficiently small, this component of $B\cap\Sigma$ is an oval with aspect ratio at most $\zeta_\Sigma/\sqrt{\varepsilon}$. Thus, $L(p)$ is contained in an oval with height 4ε and width at most $4\zeta_\Sigma\sqrt{\varepsilon}$, which implies that p has $O(k\zeta_\Sigma/\varepsilon^{1/2}) = O(kn^{1/4})$ local internal neighbors.

The ugly points lie in a neighborhood of width $\beta_\Sigma\sqrt{\varepsilon}$ around the A_3 contact curves. (We assume here that $\beta_\Sigma \gg \zeta_\Sigma$.) These curves have finite total length and (because Σ is generic) bounded curvature, so the ugly portion of the surface has area $O(\beta_\Sigma\sqrt{\varepsilon})$. The usual packing argument implies that there are $O(\beta_\Sigma/\varepsilon^{3/2}) = O(kn^{3/4})$ ugly points. □

Finally, we consider the case where p is a bad point.

Lemma 3.3. *There are $O(k^2 n\log n)$ local internal edges with bad endpoints.*

Proof: Let p be an A_3 *contact* point. To simplify our notation, let $\sigma(x) = \Sigma(x,0)$ in the local coordinate system at p. Consider a bad point $q = (x,0,\sigma(x))$, where $\beta_\Sigma\sqrt{\varepsilon} < x < \gamma_\Sigma$. Applying the Taylor expansion of σ at p, we have

$$\sigma(x) = \frac{x^2}{2r_*} + \frac{\lambda_1 x^4}{24} + \cdots,$$

where $r_* = r_*(p)$ and $\lambda_1 = \lambda_1(p)$. If γ_Σ is sufficiently small, we can approximate the principal curvatures at q by evaluating partial derivatives of Σ at $(x, 0)$.

$$\kappa_1(q) = \frac{\sigma''(x)}{(1+(\sigma'(x))^2)^{3/2}} \approx \frac{1}{r_*} + \left(\frac{\lambda_1 r_*^3 - 3}{r_*^3}\right)\frac{x^2}{2} - \cdots$$

The last approximation shows the Taylor series of $\kappa_1(q)$ with respect to x (as computed by Mathematica). A similar but much simpler calculation gives us $\kappa_2(q) \approx \kappa_2 = \kappa_2(p)$, and assuming γ_Σ is sufficiently small, we also have $r_*(q) \approx r_*$. Thus, by Equation 1, any oval $h(q, \delta)$ is an ellipse with aspect ratio

$$2\sqrt{\frac{1-\kappa_2(q)\cdot r_*(q)}{1-\kappa_1(q)\cdot r_*(q)}} \approx 2\sqrt{\frac{2r_*^2(1-\kappa_2 r_*)}{3-\lambda_1 r_*^3}}\,\frac{1}{x}$$

As in the previous proofs, let

$$\theta_\Sigma = \min_{A_3 \text{ contact points } p} \pi\sqrt{\frac{2r_*^2(1-\kappa_2 r_*)}{3-\lambda_1 r_*^3}},$$

where r, κ_2, and λ_1 are functions of the A_3 contact point p. Because Σ is generic, this is a finite positive constant. By our earlier arguments, $L(q)$ is contained in an ellipse with height 4ε and width $8\theta_\Sigma\varepsilon/x$. If we let $\deg(q)$ denote the number of local internal neighbors of q, we have $\deg(q) = O(k\theta_\Sigma/x) = O(k/x)$ by the usual packing arguments.

To compute the total number of local internal edges with bad endpoints in P, we approximate the degree sum with an integral over all bad points on the surface, multiplied by k to accommodate the oversampling allowed by our sampling condition. At the cost of another constant factor, we can evaluate this integral by considering the closest A_3 contact point to each bad point, and using the fact that the total length and maximum curvature of the A_3 curves is constant. To simplify notation, we write $f(n) \ll g(n)$ to mean $f(n) = O(g(n))$.

$$\begin{aligned}
\sum_{\text{bad } q \in P} \deg(q) &\ll kn \iint_{\text{bad } q \in \Sigma} \deg(q)\, dq^2 \\
&\ll kn \int_{A_3 \text{ points } p} \int_{\beta_\Sigma\sqrt{\varepsilon}}^{\gamma_\Sigma} \frac{k}{x}\, dx\, dp \\
&\ll k^2 n \int_{\sqrt{\varepsilon}}^{1} \frac{1}{x}\, dx = -k^2 n \ln\sqrt{\varepsilon}
\end{aligned}$$

Since $\varepsilon = \Omega(1/\sqrt{n})$, the proof is complete. □

Theorem 3.4. *There are $O(k^2 n \log n)$ local internal edges.*

3.3 Remote Internal Edges

Every medial ball touches the surface Σ at at least two points; we call each of these points is a *medial reflection* of the other(s). Almost every point p on the surface has a unique interior medial reflection, which we denote $\bar{p}$. Intuitively, the endpoints of any remote internal edge approximate a pair of medial reflections. We can formalize this intuition as follows. Let $R(p)$ denote the the *remote neighborhood* of p, the set of all possible remote internal neighbors of p.

Lemma 3.5. *Any point in $R(p)$ is in the approximate local neighborhood of the medial reflection of a point in $L(p)$.*

Proof: Let pq be a remote internal edge, and let B be a Delaunay ball for that edge. For the moment, assume that $B \cap \Sigma$ has exactly two components, one with p on its boundary and the other with q on its boundary. We can shrink B toward its center until it first becomes tangent to Σ, say at a point p' near p, and then shrink the ball toward p' until it becomes tangent to the surface again, at a point q' near q. The resulting ball is just the internal medial ball of both p' and q'. Moreover, $p' \in L(p)$, since it lies inside an almost medial ball with p on its boundary. Similarly, $q' \in L(q)$, which implies that q' is either inside or near the boundary of $L(q)$. □

If $R(p)$ has more than one connected component, we analyze each component individually as if it were the only one, by considering only parts of the surface close to that component or close to p. Intuitively, we 'move the rest of the surface out of the way'. This only increases the number of remote edges we have to count, so our final bound will be an overestimate. If ε is sufficiently small, $R(p)$ has at most three components.

Let $\bar{p}$ denote the medial reflection of p (near some component of $R(p)$). The medial reflection function μ is a diffeomorphism from a neighborhood of p to a neighborhood of $\bar{p}$. The previous lemma implies that $R(p)$ is arbitrarily close to the Minkowski sum of $\mu(N(p))$ and $N(\bar{p})$. Taylor's theorem implies that we can approximate μ by its first derivative, which is a *linear* map from the tangent plane at p to the tangent plane at $\bar{p}$.

As in the previous section, let B denote p's internal medial ball, let r denote the radius of B, let κ_1 and κ_2 denote the principal curvatures at p, and let $\bar{\kappa}_1$ and $\bar{\kappa}_2$ denote the principal curvatures at $\bar{p}$.

For the case where $\bar{p}$ is a good point, rather than delving into an exact differential analysis, we consider the following related problem: If we move the medial center c slightly, how does its closest point in the neighborhood of p (or $\bar{p}$) change? By definition, p and $\bar{p}$ are the closest points on the surface to their common medial center c. Let $B(c, \delta)$ denote a ball of infinitesimal radius δ centered on c, and let $P(p, \delta)$ denote the set of nearest neighbors to points in $B(c, \delta)$ in the neighborhood of p. Observe that $P(\bar{p}, \delta)$ contains the medial reflection of every point in $P(p, \delta)$ and vice versa.

Consider any plane π through c and p, and let σ be the intersection curve of this plane with Σ. In an small neighborhood of p, the surface normal vectors at points on σ lie close to the plane π. Thus, if we move c within the plane π, we can approximate its nearest neighbor on the surface Σ by its nearest neighbor on the curve σ. Moreover, given any point $\tilde{c}$ close to c, we can approximate its nearest neighbor on σ by projecting $\tilde{c}$ toward the center of curvature of σ at p. Thus, if κ denotes the curvature of σ at p, any point within δ of c projects to a point within $\delta/(1 - \kappa r_*)$ of p.

By definition of principal curvature, we have $\kappa_2 \leq \kappa \leq \kappa_1$. Thus, $P(p, \delta)$ contains a disk of radius $\delta/(1 - \kappa_2 r_*)$ and is contained in a disk of radius $\delta/(1 - \kappa_1 r_*)$.[4] Similarly, $P(\bar{p}, \delta)$ contains a disk of radius $\delta/(1 - \bar{\kappa}_2 r_*)$ and is contained in a disk of radius $\delta/(1 - \bar{\kappa}_1 r_*)$. We conclude that reflecting $L(p)$ across the medial axis increases its area by at most a factor of $(1 - \kappa_2 r_*)/(1 - \bar{\kappa}_1 r_*)$.

Lemma 3.6. *There are $O(k^2 n \log n)$ remote internal edges where at least one endpoint is good.*

Proof: We define yet another constant

$$\phi_\Sigma = \max_{\text{good points } \bar{p}} \left(\frac{1 - \kappa_2 r_*}{1 - \bar{\kappa}_1 r_*} \right).$$

Let p be an arbitrary point whose medial reflection $\bar{p}$ is good. Suppose the width of $N(p)$ is $4w\varepsilon$; we easily observe that $w \geq 1$. Then $\mu(N(p))$ is a convex oval of height at most $4\phi_\Sigma \varepsilon$ and width at

[4] More careful analysis implies that $P(p, \delta)$ contains an ellipse with axes of length $2\delta/(1 - \kappa_1 r_*)$ and $2\delta/(1 - \kappa_2 r_*)$ along the principal curvature directions, but this improves our results by only a constant factor.

most $4\phi_\Sigma w\varepsilon$. By Lemma 3.1, $N(\bar{p})$ is an ellipse of height 4ε and width at most $4\eta_\Sigma\varepsilon$. Thus, the area of $R(\varepsilon)$ is at most

$$16(1+\phi_\Sigma)(\eta_\Sigma+\phi_\Sigma w)\varepsilon^2 = O(w\varepsilon^2).$$

In other words, the area of $R(p)$ is at most a constant factor larger than the area of $N(p)$. Thus, the maximum possible number of remote neighbors of p is at most a constant factor more than the maximum possible number of local neighbors of p. Summing over all reflections of good points completes the proof. □

It remains to count remote edges where neither endpoint is good. If ε is sufficiently small, ugly points have only good reflections (thanks to A_3A_1 medial balls), and bad points can only have bad reflections near the same A_3 contact curve.

Lemma 3.7. *There are $O(k^2 n \log n)$ remote internal edges with two bad endpoints.*

Proof: Let p be an A_3 contact point. As in the proof of Lemma 3.2, let $\sigma(x) = \Sigma(x, 0)$ in p's local coordinate system. For small values of x, the direction of maximum curvature at $q = \sigma(x)$ lies almost exactly along the curve σ. Moreover, the medial reflection of q is arbitrarily close to the symmetric point $\sigma(-x)$. It follows that near any bad point, the medial reflection function distorts only along the direction of minimal curvature. Specifically, reflecting $L(q)$ across the medial axis increases its width, and thus its area, by a factor of $(1-\kappa_2(q)r_*(q))(1-\kappa_2(\bar{q})r^*(\bar{q}))$. Since $\kappa_2 < 1/r$ for every point on the surface, this ratio is bounded above and below by constants. Thus, the maximum number of bad remote internal neighbors of any bad point is within a constant of the number of local internal neighbors. The result now follows from Lemma 3.3. □

Theorem 3.8. *There are $O(n \log n)$ remote internal edges.*

3.4 External and Boundary Edges

Finally, to extend our analysis to external edges, we exploit the conformal invariance of Delaunay triangulations [22]. Let S^- be the largest medial ball in the interior of Σ, and let S^+ be the smallest ball containing Σ. Let ι be the unique sphere inversion that maps S^+ to S^- and vice versa. Then $\iota(\Sigma)$ is a generic smooth surface. For any pair of points $p, q \in \Sigma$, we have

$$\frac{|pq|}{\alpha_\Sigma} \le |\iota(p)\iota(q)| \le \alpha_\Sigma |pq|,$$

where α_Σ is the ratio between the radius of S^+ and the radius of S^-. Thus, $\iota(P)$ is an $(\alpha_\Sigma\varepsilon, \alpha_\Sigma^2 k)$-sample of $\iota(\Sigma)$. Thus, Theorems 3.4 and 3.8 imply that the Delaunay triangulation of $\iota(P)$ has $O(\alpha_\Sigma^4 k^2 n \log n) = O(k^2 n \log n)$ internal edges. If $\iota(p)\iota(q)$ is an edge in the Delaunay triangulation of $\iota(P)$, then pq is an edge of either the Delaunay triangulation or the anti-Delaunay triangulation of P.

At this point, we would like to argue that every external Delaunay edge in P maps to an interior Delaunay edge of $\iota(P)$, but this is not quite true. We call pq a *boundary* edge if both pq and $\iota(p)\iota(q)$ are external edges in their respective Delaunay triangulations. Let B be an empty ball with p and q on its boundary. If pq is a boundary edge, then B is centered outside Σ and $\iota(B)$ is centered outside $\iota(\Sigma)$. Fortunately, this is only possible if B is very small.

Lemma 3.9. *Every boundary edge has length at most $2\alpha_\Sigma^2\varepsilon$.*

Proof: Let pq be a boundary edge, let B be one of its Delaunay balls, and let c and r denote the center and radius of B. Since c lies outside Σ, the inverted center $\iota(c)$ lies inside $\iota(\Sigma)$. The distance from $\iota(c)$ the the boundary of $\iota(B)$ is at least r/α_Σ. Thus, $\iota(B)$ contains a point at distance at least r/α_Σ inside $\iota(\Sigma)$. Since $\iota(B)$ is a Delaunay ball, we must have $r/\alpha_\Sigma < \alpha_\Sigma\varepsilon$, or equivalently, $r < \alpha_\Sigma^2\varepsilon$. The distance between p and q is at most $2r$. □

A simple packing argument now implies that P has only $O(kn)$ boundary edges, completing the proof of the Main Theorem.

4 Discussion

4.1 Alternate Sampling Conditions

Most provably correct surface reconstruction algorithms require a sampling condition that depends on the *local feature size* of the underlying surface Σ. The local feature size of a point p, denoted $\mathrm{lfs}(p)$, is the distance from p to the closest point on the medial axis. The local feature size could be considerably smaller than the medial radius of p, but never larger. Amenta and others [1, 2, 3, 4, 12, 18, 23, 29, 30] call a set P of points in Σ an *ε-sample* if the distance from any surface point $x \in \Sigma$ to the nearest sample point in P is at most $\varepsilon\,\mathrm{lfs}(x)$.

Unfortunately, without some control of oversampling, ε-samples of any surface except the sphere can have quadratic Delaunay complexity [21]. We call p a *uniform* sample of Σ if the distance from any point $x \in \Sigma$ to the *second*-closest sample point is between $\delta\varepsilon\,\mathrm{lfs}(x)$ and $\varepsilon\,\mathrm{lfs}(x)$ for some fixed constant $0 < \delta < 1/2$ [21]. (A more general notion of *locally* uniform samples was defined by Dey *et al.* [18, 23].) Extending our Main Theorem to this sampling condition is trivial.

Theorem 4.1. *The Delaunay triangulation of any uniform ε-sample of Σ has complexity $O(n \log n)$.*

Proof: Any uniform ε-sample is an (ε', k)-sample for $\varepsilon' = \varepsilon \cdot \max_{x\in\Sigma} \mathrm{lfs}(x)$ and $k = O(\lambda_\Sigma^2)$, where

$$\lambda_\Sigma = \frac{\max_{x\in\Sigma} \mathrm{lfs}(x)}{\min_{x\in\Sigma} \mathrm{lfs}(x)}.$$

Because Σ is smooth, λ_Σ is finite and therefore constant. □

Using similar techniques, we can also extend our result to the case of random point sets. We use Poisson processes to simplify the proof, as do Golin and Na [25, 26, 27], but similar arguments involving Chernoff bounds imply the analogous result for sets of n uniformly distributed points.

Theorem 4.2. *Let P be a set of points generated by a homogeneous Poisson process with rate n over Σ. With high probability, the Delaunay triangulation of P has complexity $O(n \log^3 n)$.*

Proof: Let D be an arbitrary ε-disk, where $\varepsilon = \sqrt{(\ln n)/n}$. If n is sufficiently large, the area of D is approximately $\pi(\ln n)/n$, so the expected number of points in $D \cap P$ is approximately $\pi \ln n$. By the definition of a homogeneous Poisson process, we have $\Pr[|D \cap P| = 0] = n^{-\pi}$ and (crudely) $\Pr[|D \cap P| > 2\pi \ln n] = n^{-\Omega(\log n)}$. We can clearly cover Σ with $O(n/\log n)$ ε-disks. With probability at least $1 - O(n^{1-\pi})$, every one of these disks contains at least 1 and at most $2\pi \ln n$ points in P. Thus, with high probability, P is an $(\varepsilon, O(\log n))$-sample. □

4.2 Caveat Lector!

We must emphasize that our Main Theorem is a statement about the behavior of Delaunay triangulations *in the limit* as ε approaches zero; it says almost nothing about the complexity of Delaunay triangulations in practice. Indeed, it is quite easy to construct example where theory and practice disagree.

For example, suppose Σ is an ellipsoid with axes of length 5, 1, and $1 - 10^{-50}$. This surface is generic, so in the limit as ε approaches zero, the Delaunay complexity of any uniform ε-sample has complexity $O(n \log n)$. However, for realistic sampling densities, the surface is so close to a circular ellipsoid Σ' that moving each sample point from Σ to its nearest point on Σ' would not change the Delaunay triangulation. Thus, for realistic values of n, the worst-case complexity of the Delaunay triangulation is at least $cn^{3/2}$ for some constant c. Intuitively, the sample points only 'notice' that Σ is generic when the sample density ε is extremely small.

On the other hand, suppose Σ is the union of a generic surface (say, the Stanford bunny) with an extremely small circular cylinder (say, clipped to one ear). For reasonable values of ε, the cylinder has few sample points, so the local degenerate behavior is insignificant, and the complexity of the Delaunay triangulation will be less than $cn \log n$ for some constant c. However, in the limit as ε approaches zero, the local degenerate behavior dominates, leading to a worst-case Delaunay complexity of $\Omega(n^{3/2})$.

Naturally, we can cascade these constructions to obtain a surface whose behavior alternates between 'almost degenerate' and 'obviously generic' any finite number of times as the sampling density increases.

5 Open Problems

We conjecture that our $O(n \log n)$ upper bound is tight in the worst case; that is, for any fixed generic surface Σ, we conjecture that there is a uniform ε-sample whose Delaunay complexity is $\Omega(n \log n)$. So far, however, we have been unable to prove such a lower bound, even for simple special cases like generic ellipsoids.

The requirement of genericity appears to be necessary for our deterministic bounds, but not for the randomized bound in Theorem 4.2. We recently proved that the Delaunay triangulation of n random points on an arbitrary quadric surface—for example, a circular cylinder—has expected complexity $O(n \log n)$ [20]. We conjecture that the Delaunay triangulation of n random points on *any* fixed collection of surface patches with nonzero total area—not necessarily generic, smooth, convex, polyhedral, connected, or piecewise C^1—has complexity $O(n \operatorname{polylog} n)$ with high probability. If true, this would imply that surface samples have complex Delaunay triangulations only if both the surface and the sample points are chosen carefully.

Finally, can our results be extended to surfaces with boundaries, or with sharp ridges or points? What is required is a method to limit the interaction between smooth and discontinuous portions of the surface, similar to the recent techniques of Attali and Boissonnat [9] and Golin and Na [28] for polyhedral surfaces.

Acknowledgments. Thanks again to Herbert Edelsbrunner for asking the question that started this line of research almost three years ago. Thanks also to John Sullivan for a helpful discussion about generic surfaces.

References

[1] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete Comput. Geom.* 22(4):481–504, 1999.

[2] N. Amenta, M. Bern, and M. Kamvysselis. A new Voronoi-based surface reconstruction algorithm. *Proc. SIGGRAPH '98*, 415–412, 1998. Computer Graphics Proceedings, Annual Conference Series.

[3] N. Amenta, S. Choi, T. K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, 213–222, 2000.

[4] N. Amenta, S. Choi, and R. Kolluri. The power crust, unions of balls, and the medial axis transform. *Internat. J. Comput. Geom. Appl.* 19(2–3):127–153, 2001.

[5] E. V. Anoshkina, A. G. Belyaev, O. G. Okunev, and T. L. Kunii. Ridges and ravines: a singularity approach. *Int. J. of Shape Modeling* 1:1–11, 1994.

[6] V. I. Arnol'd. *Catastrophe Theory*, 3rd revised and expanded edition. Springer-Verlag, 1992. Translated by R. K. Thomas, after a translation by G. S. Wasserman.

[7] V. I. Arnol'd, V. V. Goryunov, O. V. Lyashko, and V. A. Vasil'ev. *Singularity Theory I.* Springer-Verlag, 1998.

[8] D. Attali and J.-D. Boissonnat. Complexity of Delaunay triangulations of points on polyhedral surfaces. Rapport de recherche 4015, INRIA Sophia-Antipolis, July 2001. ⟨http://www.inria.fr/rrrt/rr-4232.html⟩.

[9] D. Attali and J.-D. Boissonnat. A linear bound on the complexity of the Delaunay triangulations of points on polyhedral surfaces. *Proc. 7th Annu. ACM Sympos. Solid Modeling Appl.*, 139–146, 2002.

[10] I. A. Bogaevsky. Metamorphoses of singularities of minimum functions and bifurcations of shock waves of the burgers equation with vanishing viscosity. *Algebra Anal.* 1:1–16, 1989. In Russian. English translation in *Leningr. Math.* 1:807–823, 1990.

[11] I. A. Bogaevsky. Perestroikas of shocks and singularities of minimum functions. Preprint, August 2002, arXiv:math.AP/0204237.

[12] J.-D. Boissonnat and F. Cazals. Smooth surface reconstruction via natural neighbour interpolation of distance functions. *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, 223–232, 2000.

[13] J. W. Bruce and P. J. Giblin. *Curves and Singularities (2nd edition).* Cambridge University Press, New York, 1992.

[14] J. W. Bruce, P. J. Giblin, and C. G. Gibson. Symmetry sets. *Proc. Royal Soc. Edinburgh* 101A:163–186, 1985.

[15] L. N. Bryzgalova. Singularites of the maximum of a parametrically dependent function. *Funkts. Anal. Prilozh.* 11:59–60, 1977. In Russian. English translation in *Funct. Anal. Appl.* 11:49–51, 1977.

[16] H. I. Choi, S. W. Choi, and H. P. Moon. Mathematical theory of medial axis transform. *Pacific Journal of Mathematics* 181(1):57–88, 1997.

[17] S. Choi and N. Amenta. Delaunay triangulation programs on surface data. *Proc. 13th Annu. ACM-SIAM Sympos. Discrete Algorithms*, 135–136, 2002.

[18] T. K. Dey, S. Funke, and E. A. Ramos. Surface reconstruction in almost linear time under locally uniform sampling. *Abstracts 17th European Workshop Comput. Geom.*, 129–132, 2001. Freie Universität Berlin. ⟨http://www.cis.ohio-state.edu/~tamaldey/paper/recon-linear/surf.ps.gz⟩.

[19] H. Edelsbrunner. Surface reconstruction by wrapping finite point sets in space. To appear in *Discrete & Computational Geometry: The Goodman-Pollack Festschrift*, 2002. Algorithms and Combinatorics, Springer-Verlag.

[20] J. Erickson. Random points on smooth surfaces have nice Delaunay triangulations. In preparation.

[21] J. Erickson. Nice point sets can have nasty Delaunay triangulations. *Proc. 17th Annu. ACM Sympos. Comput. Geom.*, 96–105, 2001. ⟨http://www.cs.uiuc.edu/~jeffe/pubs/spread.html⟩.

[22] J. Erickson. Dense points sets have sparse Delaunay triangulations. *Proc. 13th Annu. ACM-SIAM Sympos. Discrete Algorithms*, 125–134, 2002. ⟨http://www.cs.uiuc.edu/~jeffe/pubs/screw.html⟩.

[23] S. Funke and E. A. Ramos. Smooth-surface reconstruction in near-linear time. *Proc. 15th Annu. ACM-SIAM Sympos Discrete Algorithms*, 781–790, 2002.

[24] P. J. Giblin and B. B. Kimia. A formal classification of 3D medial axis points and their local geometry. *Proc. CVPR*, I–566–I–573, 2000.

[25] M. J. Golin and H. S. Na. On the average complexity of 3D-Voronoi diagrams of random points on convex polytopes. *Proc. 12th Canad. Conf. Comput. Geom.*, 127–135, 2000. ⟨http://www.cs.unb.ca/conf/cccg/eProceedings/⟩.

[26] M. J. Golin and H. S. Na. On the average complexity of 3D-Voronoi diagrams of random points on convex polytopes. Tech. Rep. HKUST-TCSC-2001-08, Hong Kong Univ. Sci. Tech., June 2001. ⟨http://www.cs.ust.hk/tcsc/RR/2001-08.ps.gz⟩.

[27] M. J. Golin and H. S. Na. On the proofs of two lemmas describing the intersections of spheres with the boundary of a convex polytope. Tech. Rep. HKUST-TCSC-2001-09, Hong Kong Univ. Sci. Tech., July 2001. ⟨http://www.cs.ust.hk/tcsc/RR/2001-09.ps.gz⟩.

[28] M. J. Golin and H. S. Na. The probabilistic complexity of the Voronoi diagram of points on a polyhedron. *Proc. 18th Annu. ACM Sympos. Comput. Geom.*, 209–216, 2002.

[29] M. Hisada, A. G. Belyaev, and T. L. Kunii. A 3D Voronoi-based skeleton and associated surface features. *Pacific Graphics 2001*, 89–96, 2001. ⟨http://www.u-aizu.ac.jp/~belyaev/papers/pg01hbk.ps.gz⟩.

[30] H. Hiyoshi and K. Sugihara. Voronoi-based interpolation with higher continuity. *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, 242–250, 2000.

[31] V. I. Mazov. The topological classification of germs of the maximum function and minimax functions of a family of functions in generral position. *Ups. Mat. Nauk.* 37:129–130, 1982. In Russian. English translation in *Russ. Math. Surveys* 37: 127–128, 1982.

[32] V. I. Mazov. Ellipticity domains of generic families of homogeneous polynomials and extremum functions. *Funkts. Anal. Prilozh.* 19:26–36, 1985. In Russian. English translation in *Funct. Anal. Appl.* 19:102–111, 1985.

[33] E. T. Whittaker and G. N. Watson. *A Course in Modern Analysis*, 4th edition. Cambridge University Press, 1990.