Building Triangulations Using ϵ-Nets

Kenneth L. Clarkson*

March 13, 2006

Abstract

This work addresses the problem of approximating a manifold by a simplicial mesh, and the related problem of building triangulations for the purpose of piecewise-linear approximation of functions. It has long been understood that the vertices of such meshes or triangulations should be "well-distributed," or satisfy certain "sampling conditions." This work clarifies and extends some algorithms for finding such well-distributed vertices, by showing that they can be regarded as finding *ϵ-nets* or *Delone sets* in appropriate metric spaces. In some cases where such Delone properties were already understood, such as for meshes to approximate smooth manifolds that bound convex bodies, the upper and lower bound results are extended to more general manifolds; in particular, under some general conditions, the minimum Hausdorff distance for a mesh with n simplices to a d-manifold M is $\Theta((\int_M \sqrt{|\kappa(x)|}/n)^{2/d})$ as $n \to \infty$, where $\kappa(x)$ is the Gaussian curvature at point $x \in M$. We also relate these constructions to Dudley's approximation scheme for convex bodies, which can be interpreted as involving an ϵ-net in a metric space whose distance function depends on surface normals. Finally, a novel scheme is given, based on the Steinhaus transform, for scaling a metric space by a Lipschitz function to obtain a new metric. This scheme is applied to show that some algorithms for building finite element meshes and for surface reconstruction can be also be interpreted in the framework of metric space ϵ-nets.

1 Introduction

The problem considered here is that of approximating a smooth manifold by a polyhedral one. Smooth manifolds might represent, for example, the boundaries of objects and obstacles in a motion planning problem, or the graphs of functions. For many purposes, meshes (polyhedral manifolds made up of simplices) are easier to work with, and so the approximation problem arises. This problem can be attacked by first deciding where the vertices of the simplices should be, and then deciding how to connect them. The vertex placement can be done

*Bell Labs; 600 Mountain Avenue; Murray Hill, New Jersey 07974; clarkson@research.bell-labs.com

by putting the vertices on the approximated manifold, and distributing them "nicely" in an appropriate sense.

Intuitively, the niceness of the distribution of vertices would include the condition that the higher the curvature of the manifold in a given neighborhood, the more vertices in that neighborhood. This intuition can be made rigorous. Gruber[Gru04] and others have shown, for example, that for a variety of approximation requirements, the case where the manifold is the smooth boundary of a convex body can be characterized as follows: the vertices of an optimal mesh constitute a *Delone set* in a Riemannian metric D_{II} induced by the Second Fundamental Form of the manifold. This quadratic form measures the directional curvature of the manifold, so that points that are close in D_{II} can be connected by a nearly-straight curve on the surface.

Delone sets are defined more formally in Section 2 below; they are sets that are both packings and coverings: a packing has no two points too close together, and a covering S has $D(x, S)$ small for every point x in the manifold, where $D(x, S)$ is $\min_{p \in S} D(x, p)$. Metric space ϵ-nets are a particular kind of Delone set.

Riemannian metrics are defined more formally in Section 3 below; the metric D_{II} here can be described roughly as follows: the length of a very small line segment between two points on the manifold is the square root of the curvature of the manifold along that segment; the length of a path is the integral of the lengths of small segments of that path; the distance between two points is the length of the shortest path connecting the two points. We will call this the *root curvature* metric.

Gruber proved a Delone set characterization for general Riemannian manifolds, not just for D_{II}; his result has implications for optimal quantification and other problems. For example, a set S of size n that minimizes the integral of squared distances $\int_M D(x, S)^2$ must be a Delone set. In the discrete setting, such a problem is solved heuristically by the k-means (Lloyd's) algorithm; a similar heuristic can be applied in the continuous setting, as reviewed by Du *et al.*[DFG99]. The Delone requirement means that a simple and fast algorithm, the greedy method analyzed by Gonzalez [Gon85], has guaranteed approximation properties. The greedy method is discussed in Section 2, and the approximation property is Corollary 5.5.

A related line of work, starting with Eldar *et al.*[ELPZ97], applies the greedy algorithm to the problem of finding a set of vertices for piecewise-linear interpolation, in the setting of computer vision applications. Also, Peyrè and Cohen propose using the greedy algorithm together with geodesic distance measures on surfaces, for remeshing and parameterization. They consider the issue, not addressed here, of how to compute such distances efficiently; they discuss both isotropic distances and distances scaled by curvature, but not the particular measures analyzed here [PC05].

Section 5 contains a version of Gruber's proof of this Delone set condition. This new version is intended as an explication and simplification; it is also a slight generalization, to the setting of metric measure spaces. At the least, it is a bit more concise.

In the terminology of *center* and *median* problems, a cover is a good k-center. The idea of Gruber's proof, and the one here, is to show that a local search algorithm [AGK+01] for reducing $\int_M D(x,S)^2$ can make progress if S is not a Delone set. That is, if S is not a good k-center (up to a constant factor), it is not a good k-median in this particular smooth setting. A converse is also true, as shown in Corollary 5.5.

Not only has the Delone set characterization been proven for smooth convex manifolds (that is, manifolds that bound convex bodies and are smooth), but the error of approximating such manifolds has been characterized precisely: the error of best approximations that use n vertices can be given, up to a factor of $1+o(1)$, as $n \to \infty$. (The error has been found for terms of even higher order, when $d = 1$[Lud98].)

In particular, when the error measure is the Hausdorff distance, the error for a convex d-manifold M is

$$K_d(\int_M \sqrt{\kappa(x)}/n)^{2/d}(1+o(1))$$

as $n \to \infty$, where d is the dimension, K_d is a factor dependent only on the dimension, and $\kappa(x)$ is the Gaussian curvature at point $x \in M$ [Gru93]. A 2-sphere of radius r, for example, has $\int_M \sqrt{\kappa(x)} = \int_M \sqrt{1/r^2} = A/r = Cr$, for some constant C, where A is the area, so the error is within a constant factor of r/n. Note that $\int_M \kappa(x)$ is the total Gaussian curvature of the manifold, which by the Gauss-Bonet Theorem is constant for the sphere or any other convex manifold.

However, suppose the manifold is not convex, in other words, is not the boundary of a convex body. It is not hard to show, as in Section 4.1, that if the vertices are an ϵ-net with respect to a "convexified" version D_{II} of the Second Fundamental Form, then some of the same upper bounds apply as for convex manifolds, up to a constant factor. This generalizes somewhat the results of Chen *et al.*[CSXar], who showed similar upper bound results and constructions for the approximation of functions. It is shown in Section 4.2 that triangulations with small Hausdorff distance must have a number of simplices that is within a constant factor of the upper bound, under some reasonably general conditions on the triangulation and the manifold. Thus the measure $\mu_{\mathrm{II}}(M) := \int_M \sqrt{|\kappa(x)|}$, called here the "total root curvature," seems to be a fundamental measure of the difficulty of approximating a smooth manifold M.

While no such general bound on triangulation complexity in this setting seems to have appeared before, there are several results on the optimal shape of simplices in triangulations for approximating manifolds or functions. Nadler proved such a result for function approximation[Nad86], and it was extended by Heckbert and Garland[HG99]. Pottmann *et al.*[PKH+00] found the best shape for a triangle in a regular mesh (repeating that triangle over and over). (See also Shewchuk's survey of a variety of measures of the quality of triangles for function approximation, and in particular *anisotropic* measures that consider the specific function being approximated [She02].) However, it could be that the globally best triangulation uses simplices that are not optimum locally. The results of Section 4 show that, up to some conditions and a constant factor, such a situation cannot occur.

Dudley has given a bound for the approximation of convex bodies [Dud74] that has seen wide application in computational geometry. ([Cla94] is an early reference; Agarwal *et al.* survey some related work [AHPV05].) His bound is the same as the ones mentioned, with respect to dependence $1/n^{2/d}$ on the number of vertices, but is worse in its dependence on the body being approximated. However, his construction can be applied to any convex body, not just a smooth one. Moreover, while Dudley's construction gives a triangulation with the given number of vertices, the number of simplices is not similarly bounded. (A dual version bounds the number of simplices, but not vertices. Per Jeff Erickson [Eri], the number of simplices can be bounded, at the cost of losing the convexity of the approximating mesh.) For Gruber's construction, and the ones here, the number of simplices is within a constant factor (depending on dimension) of the number of vertices.

Although Dudley's construction does not explicitly compute ϵ-covers, it is clear from his proof of an error bound that the triangulation vertices it selects form an ϵ-cover of a metric space whose distance measure involves both Euclidean distance and variation in the surface normal. That is, the proof involves the construction on an input convex manifold M of a set of points S that has the following property: for any point $x \in M$, there is some point $s \in S$ such that $D_E(p, s) \leq \epsilon$, and also $D_E(v_p, v_s) \leq \epsilon$, where D_E is the Euclidean distance, v_p is the unit normal vector to M at p, and similarly for v_s. So $D(p, s) := \max\{D_E(p, s), D_E(v_p, v_s)\} \leq \epsilon$, and it's not hard to show that this function $D(., .)$ is a metric. Thus Dudley's construction involves finding an ϵ-cover in a particular metric space. The approximating polytope that results has Hausdorff distance ϵ^2 to the input manifold M.

An analogous distance measure could be defined on arbitrary manifolds, and indeed Pottmann *et al.* have done so, defining what they term a *regularized isophotic metric*, or simply *isophotic metric* [PSH$^+$04]. Section 6 shows that an ϵ-covering set for this metric is also an ϵ-covering in the root curvature metric, and therefore implies Hausdorff distance bounds for the corresponding triangulation. Approximations using this metric are also related to the Variational Shape Approximations of Cohen-Steiner *et al.* [CSAD04]; that work used a metric based on surface normals alone, and used the k-means algorithm to minimize the integral square $\int_M D(x, S)^2$ of the distance. By Gruber's results (as in Section 5), an alternative approach yielding similar results would be to use an ϵ-net in that metric.

The approximation results of Section 4 have implications for algorithmic applications. For example, suppose a motion planning problem is given, in which a plan is sought for moving an object among obstacles, where both object and obstacles have smooth boundaries. One approach to this problem might be to approximate the object and obstacles with meshes, and then use one of the many techniques for solving the problem given input with a mesh representation. If the approximations have Hausdorff distance ϵ to the original surfaces, then a solution to the approximate problem with sufficient small clearance will also be a solution to the original problem.

Except for finding the approximating meshes, the work done is proportional to functions of ϵ, and of the total root curvature $\mu_{\mathrm{II}}(M)$, where M is the collection of smooth input surfaces, and nothing else related to the input "size". That is, $\mu_{\mathrm{II}}(M)$ gives a measure of the intrinsic difficulty of polyhedral approximation to M, and so of algorithmic problems associated with M.

A similar question might be raised about nearest neighbor searching: suppose a set S of points are generated independently on a manifold, and a data structure is desired so that, given a query point q, the nearest point in S to q can be found. Are there data structures for this problem, or an approximate version of it, whose complexity depends on the total root curvature of the manifold, or other integral curvature value?

The above results generally concern manifolds without boundaries, or with simple boundaries. There is an extensive literature of algorithms for finding triangulations in two and three dimensions that have "well-shaped" simplices, and also *conform* to the complicated piecewise-linear boundaries of a region. In Section 8, we show that some algorithms for building conforming triangulations can be viewed as algorithms for finding ϵ-nets in a particular metric space.

That space, and a related one, can be described using some constructions that seem very basic, but are apparently new: given a metric space $(\mathbb{U}, D)$, and a subset $A \subset \mathbb{U}$, Section 7 shows that there is a related metric D_A on $\mathbb{U}$ such that when $D_A(x,y)$ is small,

$$D_A(x,y) \approx D(x,y)/D(x,A) \approx D(x,y)/D(y,A).$$

A related construction in that section is the following: given a 1-Lipschitz function F on $\mathbb{U}$, there is a metric D_F on $\mathbb{U}$ such that when $D_F(x,y)$ is small, it is approximately $D(x,y)/F(x)$.

These constructions are used to show that the "local feature size" used for building triangulations, and the "sampling conditions" used for proving conditions on surface reconstruction, can be used to define metrics. Under these metrics, the greedy algorithm for ϵ-net construction is similar to the Delaunay refinement algorithms of Chew and of Ruppert for building triangulations [Che89, Rup95]. The sampling conditions on surfaces[AB99] become equivalent to a Delone set condition in a corresponding metric space.

The metric space constructions are not done via a Riemannian metric, but rather via the simpler "biotope transform," and so are likely to be simpler to work with. They do not, however, give rise to length spaces. (These terms are defined in Sections 2, 3, and 7 below.)

The next two sections give some terminology and background, but before that, a little miscellaneous notation: let $D_E(a,b)$ denote the Euclidean distance between a and b; for values $\beta > 0$, x, and y, $x \leq_\beta y$ denotes the condition $x \leq (1+\beta)y$; $x \approx_\beta y$ denotes the condition that $x \leq_\beta y$ and $y \leq_\beta x$ both hold. When $x \leq_\beta y \leq_\beta z$, it follows that $x \leq (1+\beta)^2 z$, and thus $x \leq_{3\beta} z$ for $\beta < 1$. So a version of the relation holds transitively, up to a constant factor.

2 Metric Spaces, Packings, Coverings, Nets, Delone Sets

Metric spaces. Given a set $\mathbb{U}$ and distance measure $D : \mathbb{U} \times \mathbb{U} \to \Re^+$, the pair $(\mathbb{U}, D)$ is a *metric space* and D is a *metric*, if, for all $x, y, z \in \mathbb{U}$:

1. $D(x, y) = 0$ if and only if $x = y$;
2. $D(x, y) = D(y, x)$;
3. $D(x, z) \leq D(x, y) + D(y, z)$.

A space is *bounded* if $\sup_{x,y \in \mathbb{U}} D(x, y) < \infty$.

For $S \subset \mathbb{U}$ and $x \in \mathbb{U}$, the distance

$$D(x, S) := \inf_{s \in S} D(x, s);$$

this infinum will exist if S is compact, and in particular if it is finite. The notation $D(A, S) := \sup_{x \in A} D(x, S)$. (Note that this is asymmetric, and $D(A, S) \neq D(S, A)$.) The *Hausdorff distance* $H(S, A)$ between sets S and A is $\max\{D(A, S), D(S, A)\}$. Given $S \subset \mathbb{U}$, let

$$\operatorname{diam} S := \sup_{p, p' \in S} D(p, p');$$

given a collection $\mathcal{A}$ of sets, $\operatorname{diam} \mathcal{A} := \max_{\mathcal{S} \in \mathcal{A}} \operatorname{diam}(\mathcal{S})$.

Coverings, packings, Delone sets, nets. A set $S \subset \mathbb{U}$ is an:

ϵ-covering if $D(x, S) \leq \epsilon$ for all $x \in \mathbb{U}$, that is, $D(\mathbb{U}, S) = H(\mathbb{U}, S) \leq \epsilon$.

ϵ-packing if $D(s, S \setminus \{s\}) \geq \epsilon$ for all $s \in S$; that is, open balls of radius $\epsilon/2$ centered at each $s \in S$ do not meet;

(ϵ_p, ϵ_c)-Delone if S is an ϵ_p-packing and ϵ_c-covering; [1]

ϵ-net if it is (ϵ, ϵ)-Delone; that is, for any $x \in \mathbb{U}$, we have $D(x, S) \leq \epsilon$, and for any two $p, p' \in S$, we have $D(p, p') \geq \epsilon$.

Construction of nets. It is a little surprising, perhaps, that ϵ-nets even exist. However, they can be found using a simple greedy construction analyzed by Gonzalez [Gon85]: pick an element of $\mathbb{U}$ arbitrarily to be in the net E. Next repeat until E has k members, for some target size k: pick a point p in $\mathbb{U}$ whose minimum distance $D(p, E)$ is maximum in $\mathbb{U}$, and add p to E.

This simple algorithm yields an ϵ-cover, where $\epsilon := D(p, E)$, such that an $(\epsilon/2)$-cover must have at least k members. (Briefly: the output E is also an

[1] "Delone" is one transliteration of the family name of Борис Николаевич Делоне, that is, Boris Nikolaevich Delone, a Russian mathematician. Delone sets are discussed in the crystallography literature, and elsewhere. Another transliteration is "Delaunay," as in Delaunay triangulations. The constructions here will include Delaunay triangulations of Delone sets. According to J. H. Conway (via Wikipedia), Delaunay "got his surname from an Irish ancestor called Deloney, who was among the mercenaries left in Russia after the Napoleonic invasion of 1812."

ϵ-packing, by construction. If some $(\epsilon/2)$-cover E' had fewer than k members, then some two members of E would be covered by the same point $e' \in E'$, that is, be within $\epsilon/2$ of it. They would therefore be closer than ϵ to each other.) That is, the greedy algorithm is an approximation algorithm for the problem of finding the smallest ϵ such that there is an ϵ-cover with k members.

Voronoi regions. For a given metric space $(\mathbb{U}, D)$, and $S \subset \mathbb{U}$, let $\mathrm{Vor}(p, S)$ denote the Voronoi region of p in $\mathbb{U}$ with respect to S, that is, the set of $x \in \mathbb{U}$ so that p is no farther from x than any other $p' \in S$. As a formula,

$$\mathrm{Vor}(p, S) := \{x \in \mathbb{U} \mid D(x, p) = D(x, S)\}.$$

Let $C(p, S)$, the *circumradius* of $\mathrm{Vor}(p, S)$, denote

$$\sup_{x \in \mathrm{Vor}(p,S)} D(p, x),$$

that is, the maximum distance to p of points in its Voronoi region. (This is slightly abusive of terminology, because more typically the circumradius of a region A is $\inf_{p \in A} \sup_{p' \in A} D(p, p')$. Note that the circumradius of every Voronoi region of S is no more than ϵ, when S is an ϵ-cover.

Length spaces. A *length space* is a metric space for which the distance between two points is the infinum of the lengths of paths connecting the two points. As discussed just below, distances between points in Riemannian manifolds are defined using such paths, so Riemannian manifolds are length spaces.

3 Manifolds, Curvature, and Distance

This section gives a bare minimum of terminology and notation regarding manifolds and curvature. The concepts are given in most differential geometry textbooks, but not stated in the most direct way for application here.

Manifolds. A d-manifold M is a topological space that looks locally like a region of $\Re^d$; that is, there is a collection $\mathcal{V}_\mathcal{M}$ of open subsets of M, such that for each $V \in \mathcal{V}_\mathcal{M}$, there is open $U_V \subset \Re^d$ and smooth bijection $\tau_V : V \to U_V$; moreover, such *charts* τ_V and $\tau_{\hat{V}}$ must be compatible, meaning that the mapping $\tau_{\hat{V}}^{-1} \circ \tau_V$ on $V \cap \hat{V}$ must also be smooth (this holds vacuously when $V \cap \hat{V}$ is empty). The coordinates of $\tau_V(p)$ can be considered the coordinates of $p \in V$, so a chart τ_V will sometimes be called a *coordinate system* for V. There can be many different charts.

Riemannian Manifolds. A Riemannian d-manifold comprises a d-manifold M and a positive-definite quadratic form $q(x; p)$ in x, for each $p \in M$ and x in the tangent space T_p of M at p. Put another way, for each $p \in M$, there is a positive-definite matrix H_p, and $q(x; p)$ is $x^T H_p x$. Also, the entries of H_p are smooth functions of p. The form $q(x; p)$ is also called a *metric tensor*, as it can be used to define a measure of distance: the length of a curve $\gamma : [a, b] \to M$ can be given as $\int_{[a,b]} \sqrt{q(\hat{\gamma}'(t); \gamma(t))} dt$, where $\hat{\gamma}(t) = \tau_V(\gamma(t))$ when $\gamma(t) \in V \in \mathcal{V}_M$. (This expression can be extended across members of $\mathcal{V}_M$ by addition.) The distance between two points is the infinum of the lengths of the curves connecting them; that is, (M, q) is a length space. Since $\gamma'(t)$ is the tangent vector to the curve at t, $\sqrt{q(\hat{\gamma}'(t); \gamma(t))}$ can be interpreted as the length of an infinitesimal step in the direction of the curve, as measured by q. The metric tensor q also defines an associated measure of area (d-volume), $\mu_q(S) := \int_S \sqrt{\det q(\cdot; p)dp}$.

Charts of Surfaces. We will call a smooth d-manifold M embedded in $\Re^{d+1}$ a *surface*. (That is, M has *codimension* one.) Here, using smoothness and the inverse function theorem, we can assume that the charts take a particular form: for any point $p \in M$, let v_p be the unit normal to M at p, and Q_p be the $(d+1) \times d$ matrix whose columns are the unit vectors of a basis of its tangent hyperplane T_p. (We assume that the manifold has a unique unit normal at each point; this is implied by smoothness, or could almost be viewed as our definition of smoothness.) Form $(d+1) \times (d+1)$ matrix $[Q_p\ v_p]$, whose last row is v_p. The Euclidean transformation $E_p : x \rightarrow [Q_p\ v_p]^T(x - p)$ has the properties that $E_p p = 0$ and that the $x_{d+1} = 0$ hyperplane is the tangent hyperplane to $E_p M$ at p. Also, there is a neighborhood N_p of p in M such that $E_p N_p$ is a Monge patch, that is, $x_{d+1} = f(x_1, \ldots, x_d)$ for a smooth bijection f, for all $(x_1, \ldots, x_{d+1}) \in E_p N_p$. The corresponding coordinate system τ_p takes point $y \in N_p$ to $x \in \Re^d$, where $(x, f(x)) = E_p y$.

With the above conditions the Taylor expansion of f at p has its constant and linear terms equal to zero, and its quadratic term is $x^T H_p x$. (The quadric surface obtained by dropping the higher order terms in f is called the *osculating paraboloid*.)

Diagonalization. Since any Hessian H_p is symmetric, it has an eigendecomposition $H_p = S_p^T \hat{H}_p S_p$, where S_p is an orthogonal matrix of eigenvectors, and $\hat{H}_p$ is a diagonal matrix of real eigenvalues. Thus the Hessian at $0 = S_p E_p p$ of (a new version of) f is the diagonal matrix $\hat{H}_p := \operatorname{diag}(\alpha_1, \ldots, \alpha_d)$, that is, the mixed partial derivatives of f at $S_p E_p p = 0$ are all zero. With no loss of generality, we can assume that f and H have this form, and will do so from now on, so $\hat{H}_p$ will just be denoted by H_p. In this context, the above operations are called *reducing a quadratic form to a sum and difference of squares*. This is not very far from the *Morse Lemma*.

For $p \in M$, call the coordinate system constructed so far, with p at the origin, tangent plane T_p equal to the $x_{d+1} = 0$ hyperplane, and Monge patch function f with diagonal Hessian, the *coordinate system oriented to* p.

Convexification. The diagonal matrix $|H_p|$ is positive semidefinite, where $|H_p|$ has entries that are the absolute values of the entries of H_p. The matrix $H_p^c := |H_p| + \delta I$ is thus positive definite, for any $\delta > 0$, and so yields a positive definite quadratic form. Here the c means "convex", since we have replaced $x^T H_p x$ with the related convex function $x^T H_p^c x$.

The Riemannian metric tensor on M computed as $x^T H_p^c x$ at each point p (in the coordinate system oriented to p) will be denoted $q_{\mathrm{II}}(x; p)$. The "II" refers to the Second Fundamental Form (discussed below), since $x^T H_p x$ gives the value of the quadratic form II at p, for tangent vector x. This "convexification" of H_p is similar to a construction used for *anisotropic Voronoi diagrams*, and to one used by Chen *et al.* [CSXar]. The additive term δI can be made arbitrarily small, at the cost of affecting the point at which asymptotic conditions apply.

Reducing to sums and differences. Having put the Hessian at p in the diagonal form $x^T Hx = \sum_i \alpha_i x_i^2$, it is only one more step to scale x by $\sqrt{H_p^c}$, giving a region such that for $\hat{x} = \sqrt{H_p^c}x$ in it, $\hat{x}^T\hat{x} = x^T H_p^c x$. (Note that this step, unlike previous ones, is not a Euclidean transformation: it stretches along the coordinate axes.) Near $p \in M$, the metric tensor of the Euclidean distance on $\hat{x}$ corresponds to a metric tensor $q_{\mathrm{II}}(\cdot;p)$. Areas scale by $\prod_i \sqrt{\alpha_i} = \sqrt{\det H_p^c}$.

Fundamental Forms. With a coordinate system oriented to p, the matrix of the First Fundamental Form I at $0 = S_pE_pp$ is simply the $d \times d$ identity matrix I. The matrix of the Second Fundamental Form II at $0 = S_pE_pp$ is H_p, and the matrix of the Third Fundamental Form III is $H_p^2 (= H_pH_p)$.

The Riemannian metric tensor q_{I} obtained by using the First Fundamental Form gives the ordinary (inherited, natural) arc length, and the usual surface area, denoted here μ_{I}. These may also be denoted by q_E and μ_E, for Euclidean.

The metric tensor q_{III} defined by

$$q_{\mathrm{III}}(x;p) := x^T(\delta I + H_p^2)x,$$

a δ-perturbation to the Third Fundamental Form, gives an arc length which is the length of the image of the curve under the Gauss map, together with a small term equal to δ times the ordinary arc length; here the Gauss map takes a point $p \in M$ to the unit normal vector to M at p. (As with q_{II}, the δ term is added to make a positive definite quadratic form; the perturbation δI can be arbitrarily small, but not zero.) In the limit as $x \to 0$, the vector H_px approaches the difference between the unit normal $[0;-1]$ to M at 0, and the unit normal at x. Thus $x^T H_p^2 x$ approaches the squared length of that difference.[2] The corresponding area (d-volume) measure of $M' \subset M$, is

$$\mu_{\mathrm{III}}(M') := \int_{p \in M'} \sqrt{\det(\delta I + H_p^2)}dp.$$

The integrand is approximately $\det|H|$ when that quantity is not too small, that is, the absolute value of the Gaussian curvature. Up to the δ-perturbation, this integral is the area of the image of M' under the Gauss map, that is, the *total absolute curvature*.

Finally, the metric induced by the convex form of II, with a metric tensor given above as q_{II}, will be called the *root-curvature* distance, and the corresponding $\mu_{\mathrm{II}}(M') := \int_{M'} \sqrt{\det H_p^c}$ will be called the *total root curvature.* The form II at x, that is, $x^T H_p x$, is the *directional curvature* of M, in the direction x. The form q_{II} therefore is always at least as large as the absolute value of the directional curvature. Moreover, since the Taylor expansion of $f(x)$ about p is

$$f(x) = f(p) + \nabla f(p)x + x^T \nabla^2 f(p)x/2 + O(\|x\|^3),$$

[2] The (multi-dimensional) Taylor expansion of $\nabla f(x)$ at 0 is $\nabla f(x) = \nabla f(0) + \nabla^2 f(0)x + O(\|x\|^2)$, and here $\nabla^2 f(0) = H_p$ and by construction $\nabla f(0) = 0$, so $\nabla f(x) \approx H_p x$. The unit normal to M at 0 is $(0,\ldots,0,-1) \in \Re^{d+1}$, and at x is $v(x)/\|v(x)\|$, where $v(x) = (\nabla f(x), -1) \approx (H_px, -1)$. As $\|x\| \to 0$, $\|v(x)\| \to 1$, and the unit normal converges to $(H_px, -1)$. Thus the difference of the normals is H_px, up to higher-order terms, and $x^T H_p^2 x$ is its squared norm.

and $f(0) = 0$ and $\nabla f(0) = 0$ by construction for a coordinate system oriented to p, we have

$$|f(x)| = |x^T H_p x/2| + O(\|x\|^3) = q_{\mathrm{II}}(x;p)/2 + O(\|x\|^3).$$

For p and p' close enough together, $D_{\mathrm{II}}(p,p')^2 \approx q_{\mathrm{II}}(p'-p;p)$, and so $D_{\mathrm{II}}(p,p')^2$ is thus an upper bound on $|f(p')|$, and again, $f(p')$ is the distance of $p' \in M$ to the tangent plane at p, the deviation of M from flatness between p and p'. This is discussed more formally as Lemma 4.1.

In Section 6, a metric implicitly used by Dudley for convex boundaries, and called *isophotic* by Pottmann *et al.* will be discussed; its metric tensor is $q_{\mathrm{I+III}} := q_{\mathrm{I}} + q_{\mathrm{III}}$. Since I is positive definite, and III is positive semidefinite, no "convexification" is needed for q_{I} or $q_{\mathrm{I+III}}$, and only the δ perturbation is needed to make q_{III} positive definite.

The distances on M implied by q_X, for $X \in \{\mathrm{I}, \mathrm{II}, \mathrm{III}, \mathrm{I+III}\}$, will be denoted by the corresponding D_X. Combined with the corresponding area measures μ_X, the *metric measure spaces* (M, D_X, μ_X) are obtained. A ball with center x and radius ϵ in metric D_X will be denoted $B_X(x,\epsilon)$.

As discussed further in Section 5, a metric measure space (M, D, μ) will be called *dimension regular* if there is a number d and constant C such that for any $x \in M$, $\mu(B(x,\epsilon)) \approx_C \epsilon^d$, for sufficiently small ϵ. We will generally assume that a manifold M is such that (M, D_X, μ_X) is dimension regular.

The following lemma is straightforward.

Lemma 3.1 *If the compact metric measure space* $(\mathbb{U}, D, \mu)$ *is dimension regular with dimension* d*, then the size of an* ϵ*-net of* $\mathbb{U}$ *is* $\Theta(\mu(\mathbb{U})/\epsilon^d)$ *as* $\epsilon \to 0$*.*

Gruber[Gru93], Lemma 1, proves such a bound for Riemannian manifolds that is tight up to lower order terms.

4 Hausdorff Approximations

The D_{II} distance will be the main concern here; the next lemma gives the basic relations between D_{II} distance and deviation from linearity. The lemma follows from the lemmas of Gruber[Gru93, Gru04], but a proof is included for completeness.

Lemma 4.1 *For any compact smooth* d*-surface* M*, value* λ *with* $0 < \lambda < 1$*, and point* $r \in M$*, there is neighborhood* $V_r(\lambda)$ *of* r *such that for* $p, p' \in V_r(\lambda)$*,*

$$D_E(p, T_{p'}) \leq_\lambda D_{\mathrm{II}}(p,p')^2 \approx_\lambda \|\sqrt{H_r^c}(\tau_r(p-p'))\|^2$$

in the coordinate system oriented to r*, where* $D_E(p, T_{p'})$ *is the minimum Euclidean distance from* p *to the tangent plane at* p'*, and* $\tau_r p$ *is the orthogonal projection of* p *onto* T_r*. Also, for* $\hat{V} \subset V_r(\lambda)$*,*

$$\mu_{\mathrm{II}}(\hat{V}) \approx_\lambda \mu(\tau_r \hat{V})\sqrt{|\det H_r^c|}.$$

Proof: By the smoothness of M, there is a neighborhood V of r where for all $p \in V$ and all $x \in \Re^{d+1}$, in a coordinate system oriented to r,

$$
\begin{aligned}
x^T \nabla^2 f(p) x \leq_\lambda q_{\mathrm{II}}(x;p) &\approx_\lambda q_{\mathrm{II}}(\tau_r x; r) \\
&= (\tau_r x)^T H_r^c (\tau_r x) = \|\sqrt{H_r^c}(\tau_r x)\|_2^2.
\end{aligned}
$$

(Here, again, $f(x)$ is the Monge function for the manifold in the neighborhood, and $\tau_r x$ is the projection of x to $\Re^d$, that is, setting the last coordinate to zero.) For the dilated coordinates $y := \sqrt{H_r^c}\tau_r x$, then, $q_{\mathrm{II}}(x;p) \approx_\lambda y^T y$. The latter is the metric tensor of Euclidean distance, so $D_{\mathrm{II}}(p,p')^2 \approx_\lambda \|\sqrt{H_r^c}(\tau_r(p-p'))\|^2$, as claimed.

The tangent $T_{p'}$ is the best linear approximation to f at p'. By Taylor's theorem with Lagrange remainder,

$$
f(p) = f(p') + (p-p')^T \nabla f(p') + (p-p')^T \nabla^2 f(p^*)(p-p'),
$$

where p^* is on the line segment from p to p'. So

$$
\begin{aligned}
D_E(p, T_{p'}) &\leq |f(p) - (f(p') + (p-p')^T \nabla f(p'))| \\
&= |(p-p')^T \nabla^2 f(p^*)(p-p')| \\
&\leq_\lambda (p-p')^T H_{p^*}^c (p-p') \\
&\approx_\lambda (p-p')^T H_r^c (p-p')| \\
&= \|\sqrt{H_r^c}(\tau_r(p-p'))\|^2 \approx_\lambda D_{\mathrm{II}}(p,p')^2.
\end{aligned}
$$

Using the properties of the $\approx_\lambda$ relation, and renaming λ completes the proof of the first relation of the lemma.

The area relation comes from the change of variable theorem of calculus, but can also be described as follows. For $p' \in B_{\mathrm{II}}(p,\epsilon) \subset V_r(\lambda)$, the relation

$$
D_{\mathrm{II}}(p,p')^2 \approx_\lambda \|\sqrt{H_r^c}(\tau_r(p-p'))\|^2
$$

implies that $\sqrt{H_r^c}\tau_r B_{\mathrm{II}}(p,\epsilon)$ is contained in

$$
B_E(\sqrt{H_r^c}\tau_r p, \epsilon(1+\lambda)),
$$

and contains the concentric ball with radius $\epsilon/(1+\lambda)$, so on one side,

$$
\mu(\sqrt{H_r^c}\tau_r B_{\mathrm{II}}(p,\epsilon)) \leq \epsilon^d (1+\lambda)^d \leq \epsilon^d(1+2d\lambda),
$$

for small enough λ, and this measure is greater than $\epsilon^d/(1+2d\lambda)$, so

$$
\mu(\sqrt{H_r^c}\tau_r B_{\mathrm{II}}(p,\epsilon)) = \sqrt{\det H_c^r}\mu(\tau_r B_{\mathrm{II}}(p,\epsilon)) \approx_{2d\lambda} \epsilon^d.
$$

Since $\mu_{\mathrm{II}}(B_{\mathrm{II}}(p,\epsilon)) = \epsilon^d$, the area relation follows for all sufficiently small balls, and so for larger regions, after redefining λ appropriately. ■

4.1 Upper Bound

The upper bound construction uses an ϵ-net on the surface, in the D_{II} metric, and then computes a Delaunay triangulation in the D_{II} metric. The mesh $\mathcal{T}$ then has the same vertex set, and for each D_{II}-Delaunay simplex, a Euclidean simplex with the same vertices.

To use this construction, we need to ensure that the ϵ-net E has a well-defined D_{II}-Delaunay triangulation. The results of Leibon and Letscher [LL00] imply this, for ϵ small enough. The key property needed is that for $d+1$ points close enough together, there is exactly one circumscribing sphere. The Delaunay triangulation is then implied, as usual, by the set of circumspheres of points in the ϵ-net E that are empty. Here empty means that the open ball bounded by such a sphere contains no points of E.

The simplices of the triangulation, combinatorially, are the sets of points of E that determine an empty circumsphere. It will be helpful to consider also the simplices as a geometric subdivision of M. Such simplices can then be defined as the cells of the power diagram of the triangulation circumcenters. The power diagram is a kind of weighted Voronoi diagram; here the weight of a circumcenter is picked to be the radius of its circumsphere. (See [ACK01], for example, or an allusion in [LL00].) That is, each empty circumsphere with center c has radius r_c, and the region (Delaunay "simplex") of c is

$$\{p \in M \mid D_{\mathrm{II}}(c,p)^2 - r_c^2 \le \min_{\text{center } c'} D_{\mathrm{II}}(c',p)^2 - r_{c'}^2\}$$

By definition, every point of M is in some face of the triangulation. It is not hard to show that every point of M will be in some circumsphere, and that each point of E determining the circumsphere of c is in the simplex of c, and also that the simplex of c is contained in the circumsphere of c. Also, since E is an ϵ-net, Delaunay neighbors are no more than 2ϵ apart: otherwise, the center of the circumsphere of the neighbors is more than ϵ from them, and so more than ϵ from any point in E.

It seems likely that the full power of [LL00] should not really be needed, because the constructed Delaunay triangulation will be equivalent to the Delaunay triangulation based on a metric only slightly distorted from Euclidean. From the previous lemma, each neighborhood $V_r(\lambda)$ can be projected and scaled as $\sqrt{H_r^c}\tau_r V_r(\lambda)$ (in the coordinate system oriented to r), so that Euclidean distance in the transformed space is approximately equal to the D_{II} distance between the corresponding points in $V_r(\lambda)$. It follows that the anisotropic Delaunay triangulation based on D_{II} will look, at the local scale of interest, very much like an ordinary Euclidean Delaunay triangulation in $\sqrt{H_r^c}\tau_r V_r(\lambda)$.

Note that it is not claimed that the mesh $\mathcal{T}$ so constructed is well-behaved in every way, for example, with respect to orientation-reversal.

An alternative approach may be to find triangulations in each patch $V_p(\lambda)$ by using the lower convex hull of the points, as lifted to a quadratic approximation surface. To ensure consistency of the triangulation across domains, a blending of quadratic approximations using partitions of unity could be done.

Theorem 4.2 *A compact smooth d-surface M has a triangulation $\mathcal{T}$ comprising m vertices and $O(m)$ simplices, and with Hausdorff distance*

$$H(M, \mathcal{T}) = O((\mu_{\mathrm{II}}(M)/m)^{2/d})$$

as $m \to \infty$, where the constant factors in the asymptotic bounds depend only on the dimension.

Proof: As discussed above, pick an ϵ-net E in the D_{II} metric as the vertex set of the triangulation. By Lemma 3.1, with m points, such a net will have $\epsilon \leq K_d(\mu_{\mathrm{II}}(M)/m)^{1/d}$.

By the Lebesque Number Lemma and compactness of M, there is a $\gamma > 0$ such that for every $p \in M$, the ball $B_{\mathrm{II}}(p, \gamma)$ is contained in some member of the open cover $\{V_r(\lambda) \mid r \in M\}$. Choose m large enough that $\epsilon \leq \gamma/5$, so that every ball $B_{\mathrm{II}}(p, 5\epsilon) \subset V_r(\lambda)$, for some r, where $p \in E$.

As discussed above, the interpolating mesh $\mathcal{T}$ will have the same vertices as the D_{II}-Delaunay triangulation $\mathcal{T}_{\mathrm{II}}$, and for each simplex t_{II} in $\mathcal{T}_{\mathrm{II}}$, $\mathcal{T}$ will have a Euclidean simplex t with the same vertex set.

Consider a point $p' \in B_{\mathrm{II}}(p, \epsilon)$. There is a simplex $t \in \mathcal{T}$, with vertices in $B_{\mathrm{II}}(p, 4\epsilon)$, such that t is above p', that is, $\tau_r p' \in \tau_r t$. This can be proven as follows. Consider the set $\mathcal{T}_{\mathrm{II}}(p, 3\epsilon)$ of all t_{II} that meet $B_{\mathrm{II}}(p, 3\epsilon)$. From the above discussion, such simplices are contained in $B_{\mathrm{II}}(p, 5\epsilon)$. Let $\mathcal{T}(p, 3\epsilon)$ denote the corresponding simplices of $\mathcal{T}$. Since every point of $V_r(\lambda)$ is in some simplex of $\mathcal{T}_{\mathrm{II}}$, $\mathcal{T}_{\mathrm{II}}(p, 3\epsilon)$ covers $B_{\mathrm{II}}(p, 3\epsilon)$. There is a continuous mapping from each t_{II} onto the corresponding t, and so a continuous mapping from $B_{\mathrm{II}}(p, 3\epsilon)$ onto the union of the simplices of $\mathcal{T}(p, 3\epsilon)$. Hence the union of $\mathcal{T}(p, 3\epsilon)$ has no holes, and moreover, since the boundary of $\mathcal{T}_{\mathrm{II}}(p, 3\epsilon)$ is outside $B_{\mathrm{II}}(p, 3\epsilon)$, the boundary of $\tau_r \mathcal{T}(p, 3\epsilon)$ must be outside $\tau_r B_{\mathrm{II}}(p, \epsilon)$, for λ small enough. So $\tau_r p'$ is in some $\tau_r t$, for $t \in \mathcal{T}$ with $\tau_r t \subset \tau_r B_{\mathrm{II}}(p, 4\epsilon)$.

By the previous lemma, every vertex of t is within Euclidean distance $16\epsilon^2(1+\lambda)$ of T_p, and so some point of t is within that Euclidean distance of $\tau_r p'$. Since p' is no more than $\epsilon^2(1+\lambda)$ from $\tau_r p'$, it follows that some point of t is within $17\epsilon^2(1+\lambda)$ of p'.

Since for every point $p' \in M$, there is some $p \in E$ within D_{II} distance ϵ, it follows that $D_E(M, \mathcal{T}) \leq 17\epsilon^2(1+\lambda)$.

A similar, but simpler argument shows that $D_E(\mathcal{T}, M) \leq_C \epsilon^2$: for $t \in \mathcal{T}$, pick a vertex p of T; since no point of t_{II} is farther than 2ϵ from p, no point of t is more than $4\epsilon^2(1+\lambda)$ from T_p. For each point in $\tau_r t$, there is point of M with the same projection, and also within $4\epsilon^2(1+\lambda)$ of T_p. Therefore any point of t is within $8\epsilon^2(1+\lambda)$ of some point of M.

The theorem follows. ∎

4.2 Lower Bound

When the diagonal entries of the diagonal matrix H (as discussed in Section 3) all have the same sign, the second fundamental form is convex (or concave), and so the first approximate inequality of Lemma 4.1 becomes an approximate equality, and the upper bound of the above theorem can become a lower bound as well. However, when the signs are mixed, no such direct relationship between $D_{\rm II}$ and approximation error is possible, because in the mixed case, the manifold (and some of it tangents) may contain straight line segments. The distance $D_{\rm II}$ thus does not allow lower bounds. As will be shown, though, the measure $\mu_{\rm II}$ does, for Hausdorff distance: roughly, if a simplex is large in measure, its error must be also. After a lemma, this is shown for any triangulation of a pure-quadric patch, and then for somewhat-restricted triangulations of somewhat-restricted manifolds.

4.2.1 Lower Bound for Function Interpolation

The next lemma gives the basic relation between interpolation error and function value, for a quadratic function. Recall the form of the function assumed in the lemma is no loss of generality, except for the constraints on the α_i's.

Lemma 4.3 *For $x := (x_1, \ldots, x_d) \in \Re^d$ let*

$$f(x) := \sum_i \alpha_i x_i^2 = x^T D x,$$

where $\alpha_i = \pm 1$ for all i, and $D := \mathrm{diag}(\alpha_1, \alpha_2, \ldots, \alpha_d)$. Then the maximum error in linearly interpolating f between $x \in \Re^d$ and $y \in \Re^d$ is $|f(x-y)|/4$.

Proof: Noting that $x^T D y = y^T D x$, the error of interpolation at a point β of the way from x to y is the absolute value of

$$\begin{aligned}
&f(x)+\beta(f(y)-f(x))-f(x+\beta(y-x))\\
&\quad = x^T D x + \beta(y^T D y - x^T D x) - (x+\beta(y-x))^T D(x+\beta(y-x))\\
&\quad = x^T D x + \beta(y^T D y - x^T D x) - x^T D x - \beta(y-x)^T D x - \beta x^T D(y-x) - \beta^2 (y-x)^T D(y-x)\\
&\quad = \beta(y^T D y - x^T D x - (y-x)^T D x - x^T D(y-x)) - \beta^2 (y-x)^T D(y-x)\\
&\quad = (\beta-\beta^2)(y-x)^T D(y-x)\\
&\quad = \beta(1-\beta) f(y-x)
\end{aligned}$$

Since $\beta(1-\beta)|f(y-x)|$ is maximum at $\beta = 1/2$, the lemma follows. ■

Lemma 4.4 *For $x := (x_1, \ldots, x_d) \in \Re^d$ let*

$$f(x) := \sum_i \alpha_i x_i^2 = x^T D x,$$

where $\alpha_i = \pm 1$ for all i, and $D := \mathrm{diag}(\alpha_1, \alpha_2, \ldots, \alpha_d)$. Then for a simplex $t \subset \Re^d$, if the maximum error of linearly interpolating f within t is no more than ϵ, then the volume of t is no more than $\epsilon^{d/2} 2.5^d/\sqrt{d!}$.

Here the linear interpolation is assumed to assign $f(v)$ to each vertex v of t, that is, the error is zero at the vertices.

Proof: We can assume that the coordinates for which $\alpha_i = 1$ are the first d_+, for some $d_+ \leq d$, and the coordinates with $\alpha_i = -1$ are the following $d_- := d - d_+$ ones.

Since $x - y = (x - z) - (y - z)$, and error estimates will be by way of the previous lemma, we can assume that one vertex of t is the origin. Let $[a_j\ b_j]$, for $j = 1, \ldots, d$, be the other vertices, where a_j has d_+ coordinates, and b_j has d_- coordinates. Then the previous lemma, and assuming the error bound of ϵ, imply that

$$|f([a_j\ b_j] - [0\ 0])| \leq 4\epsilon,$$

for $j = 1, \ldots, d$. For j and k in $1, \ldots, d$,

$$\begin{aligned}
4\epsilon &\geq |f([a_j\ b_j] - [a_k\ b_k])| \\
&= |(a_j - a_k)^2 - (b_j - b_k)^2| \\
&= |a_j^2 - 2a_j \cdot a_k + a_k^2 - (b_j^2 - 2b_j \cdot b_k + b_k^2)| \\
&= |f([a_j\ b_j]) + f([a_k\ b_k]) - 2(a_j \cdot a_k - b_j \cdot b_k)|,
\end{aligned}$$

and so $|a_j \cdot a_k - b_j \cdot b_k| \leq 6\epsilon$. If we define matrix A as having rows a_j, and B as having rows b_j, then the above says that $\|AA^T - BB^T\|_\infty \leq 6\epsilon$, where the matrix norm is simply the maximum of the absolute values of the entries. We can also write $AA^T - BB^T$ as $[A\ B][A\ {-B}]^T$, where $[A\ B]$ is a $d \times d$ matrix, so we have

$$\|[A\ B][A\ {-B}]^T\|_\infty \leq 6\epsilon.$$

As is commonly known, the volume of t is $|\det([A\ \ B])|/d!$ in this notation. Therefore, using standard facts about the determinant, including that $\det X \leq d!$, if $\|X\|_\infty \leq 1$, and $\det(\epsilon X) = \epsilon^d \det(X)$ for $d \times d$ matrix X,

$$\begin{aligned}
d! \operatorname{Vol}(t) &= |\det([A\ B])| \\
&= \sqrt{|\det([A\ B])||\det([A\ {-B}]^T)|} \\
&= \sqrt{|\det([A\ B][A\ {-B}]^T)|} \\
&\leq (6\epsilon)^{d/2}\sqrt{d!} \\
&\leq \epsilon^{d/2} 2.5^d \sqrt{d!},
\end{aligned}$$

and the lemma follows. ■

4.2.2 Lower Bound for More General Manifolds

We can generalize this lower bound to a broader setting, but so far a completely general statement for all smooth manifolds and all triangulations has been elusive. A few conditions on the triangulation $\mathcal{T}$ and manifold M are needed. We start with the assumption that the vertices of the approximating mesh are on the manifold.

Small Diameter. The property of smoothness allows us to consider triangulations of small neighborhoods, but it is difficult to handle simplices that span many such neighborhoods, even though the lower bound lemma above implies that they must have small measure in any particular neighborhood, if the error is small. We will simply consider triangulations with simplices each of which has small diameter in the D_{II} metric. Since only the vertices of a simplex are assumed to be on the surface M, we need a little more: let $\mathrm{Vor}(t, M)$ denote the Voronoi region of simplex t in M, the set of $p \in M$ for which some point of t is closest in the triangulation $\mathcal{T}$, in Euclidean distance. We will assume that $\max_{t\in\mathcal{T}} \mathrm{diam}_{\mathrm{II}}(\mathrm{Vor}(t, M))$ is small enough, where the threshold depends on M.

Locally Roughly Quadratic. We also need, to apply the above lemma, a patch of the manifold to behave "enough" like a quadratic function. Specifically, we need the condition that if the interpolation error of a local Monge patch is small, then the error of its local quadratic approximation is also small.

For $p \in M$ and $\lambda > 0$, pick $\gamma > 0$ such that $B_{\mathrm{II}}(p, \gamma) \subset V_p(\lambda)$ (where $V_p(\lambda)$ was defined in Lemma 4.1), and let

$$L_\gamma(p) := \sup_{\substack{p',p''\in B_{\mathrm{II}}(p,\gamma)\\ y=\tau_p p', x=\tau_p p''-y}} \left|\frac{x^T H_p x}{f(y + x/2) - (f(y) + f(y + x))/2}\right|, \qquad (1)$$

where as usual, $f(p)$ is the Monge patch function in the coordinate system oriented to p, and $H_p = \nabla^2 f(p)$. If the denominator is zero, the fraction is taken as infinite if $x^T H_p x$ is nonzero, and as one otherwise. Also, it will be convenient to make $L_\gamma(p)$ infinite if H_p has determinant zero. Say that M if *locally roughly quadratic* if there is some $\gamma > 0$, $\beta > 0$, and $\psi > 0$ such that

$$\mu_{\mathrm{II}}(\{p \in M \mid L_\gamma(p) \le \beta\}) > \psi\mu_{\mathrm{II}}(M),$$

that is, L_γ is small over most of M.

While this definition was constructed to fill the needs of the proof, the denominator $f(y + x/2) - (f(y) + f(y + x))/2$ is proportional to $x^T H_p x$ when $f(x) = x^T H_p x$, and so the $L_\gamma(p)$ function is large only when higher order terms in the Taylor expansion of f are large; that is, when f varies too much from quadratic. In particular, if x is such that $x^T H_p x = 0$, then f must be linear in the x direction in the relevant neighborhood of p.

We will assume that the values γ, β, and ψ in this definition are fixed, but unspecified. They affect the constant factor in bound below.

Local Average Aspect Ratio. Given a triangulation $\mathcal{T}$ and a ball $B := B_{\mathrm{II}}(e, \delta)$ in some $M' \subset M$, let $N(B)$ denote the set of simplices in $\mathcal{T}$ that contain points that are nearest neighbors of points in B. That is, each simplex in $N(B)$ has a point p such that for some $p' \in B$, $D_{\mathrm{II}}(p, p') = D_{\mathrm{II}}(p', \mathcal{T})$. Consider the coordinate system oriented to e, so that $\tau_e(N(B))$ is the projection of the simplices in $N(B)$ onto the tangent T_e at e.

Define the *local average aspect ratio* of $\mathcal{T}$ as

$$\mathrm{LAAR}(\mathcal{T}; M', \delta) := \sup_{e\in M', B=B_{\mathrm{II}}(e,\delta)} \frac{\sum_{t\in N(B)} \mathrm{SA}_E(\tau_e(t))}{\sum_{t\in N(B)} \mu_E(\tau_e(t))},$$

where SA() is the (Euclidean) surface area. This quantity will be required to be bounded above.

A key property here is that for a convex body P and unit ball B,

$$\mu_E(P + \epsilon B) = \mu_E(P) + \epsilon\, \mathrm{SA}(P) + O(\epsilon^2) \tag{2}$$

as $\epsilon \to 0$.

Of course, it is sufficient for this condition to hold that a similar one hold for every simplex $t \in \mathcal{T}$ individually, but the weaker local average condition is enough for the lower bound. Also, suppose each simplex $t \in \mathcal{T}$ has a containing hyperplane that is nearly parallel to the tangent hyperplane to M at points in $\mathrm{Vor}(t)$. Here the aspect ratio of t, in this surface-area-to-volume sense, is about the same as its projected version. So it is sufficient that every simplex to have small aspect ratio, and to have a normal vector that has small angle to the normal at the manifold region it is approximating.

Theorem 4.5 *Suppose M is a d-surface which is roughly locally quadratic, as defined just above, for some γ, β, and ψ. Then there is a value γ' depending on M, and a constant K_d depending only on the dimension, so that the following hold. Suppose $\mathcal{T}$ is a Euclidean triangulation near to M, whose vertex set is a subset of M, and such that* $\max_{t \in \mathcal{T}} \mathrm{diam}_{\mathrm{II}}(\mathrm{Vor}(t, M)) \le \gamma'$, *and*

$$\mathrm{LAAR}(\mathcal{T}; M, \gamma') \le 1/H(\mathcal{T}, M).$$

Then when the number of simplices $|\mathcal{T}|$ is large enough,

$$H(\mathcal{T}, M) \ge \frac{K_d}{\beta}(\psi \mu_{\mathrm{II}}(M)/|\mathcal{T}|)^{2/d},$$

for some K_d depending only on the dimension.

Proof:

The proof has three major steps, that show the following:

1. For γ' small enough, every $t \in \mathcal{T}$ belongs to a "well-behaved" patch of M, using the diameter bound;

2. There is some $t^* \in \mathcal{T}$ that is "large", using the small local average aspect ratio;

3. The error of linear interpolation (vertical distance) of f on t^* is large, using the hypothesis that M is roughly locally quadratic, and Lemma 4.4.

Consider the compact set

$$M_{\beta,\gamma} := \{p \in M \mid L_\gamma(p) \le \beta\},$$

and its open cover

$$\{V_p(\lambda) \cap M_{\beta,\gamma} \mid p \in M\}.$$

Here $V_p(\lambda)$ is again as in Lemma 4.1, and $\lambda > 0$ is small enough, certainly less than one tenth. Applying the Lebesque Number Lemma, there is some radius γ' so that all balls in $M_{\beta,\gamma}$ with radius no more than $3\gamma'$ are contained in some member of that open cover. Set $\gamma' := \min\{\gamma', \gamma\}$, and find a γ'-net E of $M_{\beta,\gamma}$.

For a simplex $t \in \mathcal{T}$, let $\mathrm{Vor}(t)$ denote its Voronoi region in $M_{\beta,\gamma}$.

Since E is a γ'-net, for any point $x \in \mathrm{Vor}(t)$, there is some $e \in E$ within D_{II} distance γ' of x. For such x and e, the assumption $\mathrm{diam}_{\mathrm{II}}(\mathrm{Vor}(t)) \le \gamma'$ implies that $\mathrm{Vor}(t) \subset V_e(\lambda) \cap B_{\mathrm{II}}(e, 2\gamma')$, and also that $L_\gamma(x) \le \beta$, as defined by (1).

Thus every simplex $t \in \mathcal{T}$ has $\mathrm{Vor}(t)$ contained in some well-behaved $B(e, 2\gamma')$, concluding part one of the proof outline.

Consider now the coordinate system oriented to e, as discussed in §3. Since $\mathrm{Vor}(t) \subset B_{\mathrm{II}}(e, 2\gamma')$, each vertex v of t is in $B_{\mathrm{II}}(e, 2\gamma')$. Hence $\tau_e(v) \in \tau_e(B_{\mathrm{II}}(e, 2\gamma'))$, and so $\tau_e(t) \subset \tau_e(B_{\mathrm{II}}(e, 3\gamma'))$, for small enough λ. (Here the factor of three in the radius allows for nonconvexity of $\tau_e(B_{\mathrm{II}}(e, 2\gamma'))$ due to variation in $q_{\mathrm{II}}(;p)$ over $V_p(\lambda)$.)

For given $e \in E$ and $B := B_{\mathrm{II}}(e, \gamma') \cap M_{\beta,\gamma}$, let $N(B)$, as above, be the set of $t \in \mathcal{T}$ such that $B \cap \mathrm{Vor}(t)$ is not empty. By the assumption regarding the local average aspect ratio,

$$\sum_{t \in N(B)} \mu_E(\tau_e(t)) \ge H(\mathcal{T}, M) \sum_{t \in N(B)} \mathrm{SA}_E(\tau_e(t)).$$

So by (2) and Lemma 4.1,

$$\begin{aligned}
2 \sum_{t \in N(B)} \mu_E(\tau_e(t)) &\ge \sum_{t \in N(B)} \mu_E(\tau_e(t)) + H(\mathcal{T}, M)\,\mathrm{SA}_E(\tau_e(t)) \\
&\ge \sum_{t \in N(B)} \mu_E(\tau_e(t) + B_E(0, H(\mathcal{T}, M)))/2 \\
&\ge \sum_{t \in N(B)} \mu_E(\tau_e(\mathrm{Vor}(t)))/2 \\
&\ge \mu_{\mathrm{II}}(B)/2(1+\lambda)\sqrt{|\det H_e|},
\end{aligned}$$

for small enough $H(\mathcal{T}, M)$. (Here $B_E(0, H(\mathcal{T}, M))$, as with $\tau_e(t)$, is in the tangent hyperplane to M at e.)

Since E is a γ'-net, $\mathrm{Vor}(t)$ cannot be contained in too many balls $B(e, 2\gamma')$, at most $\hat{K}_d$ for a constant $\hat{K}_d$ at most exponential in d. Letting $e(t) \in E$ denote a member of E such that $t \in N(B_{\mathrm{II}}(e, \gamma'))$, and $h_e := \sqrt{|\det H_e|}$, we have, using the definition of "roughly locally quadratic,"

$$\sum_{t \in \mathcal{T}} h_{e(t)} \mu_E(\tau_{e(t)}(t)) \ge \sum_{e \in E} \mu_{\mathrm{II}}(B)/4\hat{K}_d \ge \psi \mu_{\mathrm{II}}(\mathbb{U})/4\hat{K}_d.$$

Therefore, there is some $t^* \in \mathcal{T}$ such that

$$h_{e(t^*)} \mu_E(\tau_{e(t^*)}(t^*)) \ge K_d \psi \mu_{\mathrm{II}}(\mathbb{U})/|\mathcal{T}|,$$

where $K_d := 1/4\hat{K}_d$.

Thus there is some "large" $t^* \in \mathcal{T}$, concluding part two of the proof outline.

Applying the transformation $\hat{x} = \sqrt{|H_e|}x$, where $e := e(t^*)$, we obtain a simplex $\hat{t} = \sqrt{|H_e|}\tau_e(t^*)$ in the transformed coordinates with

$$\mu_E(\hat{t}) \geq K_d \psi \mu_{\mathrm{II}}(\mathbb{U})/|\mathcal{T}|. \tag{3}$$

Moreover, the function $x^T H_e x$ is equal to $\hat{x}^T D \hat{x}$, where D is a diagonal matrix with entries equal to ± 1. Applying Lemma 4.4, we have that the error (vertical distance) of linear interpolation within t^* is at least $K_d(\psi\mu_{\mathrm{II}}(\mathbb{U})/|\mathcal{T}|)^{2/d}$, using a new value of K_d. Applying the hypothesis that M is "locally roughly quadratic," the error of linear interpolation of f is at least as large at that for $x^T H_e x$, divided by β, and so is

$$\frac{K_d}{\beta}(\psi\mu_{\mathrm{II}}(\mathbb{U})/|\mathcal{T}|)^{2/d},$$

the bound of the theorem statement.

We are interested in the Hausdorff distance from t to M, and not linear interpolation of f; however, the Hausdorff distance includes the distance from each point of t to M, and by construction λ bounds the angle between the "vertical" normal at e and the unit normal at the nearest neighbor in M to a point in t. The theorem follows, after again adjusting and renaming K_d. ■

5 Optimal Quantization Sets are Delone

It may be of interest to judge approximations using other distance measures, for example, the average distance from the mesh to the manifold, rather than the maximum distance. Such formulations lead to the *optimal quantization problem* of information theory.

For a metric measure space $(\mathbb{U}, D, \mu)$ and penalty function g, the optimal quantization problem is to find a set $S \subset \mathbb{U}$ such that $\int g(D(x,S))d\mu(x)$ is as small as possible. Gruber[Gru04] found tight bounds for the optimal quantization problem, when $\mathbb{U}$ is a Riemannian manifold, and D and μ are the associated metric and measure, respectively. The paper also shows that sets S that are optimal quantizers (solve the above problem) are also Delone sets, and that characterization is said to be "the hardest part of the proof." This section gives a proof of this characterization, using the natural assumptions of "Voronoi regularity" and "dimension regularity". For Riemannian manifolds these conditions naturally follow from some of Gruber's preliminary lemmas. The proof given here applies to the penalty function $g(z) = z^2$, but the extension to other exponents is trivial; Gruber has shown that the results apply to an even broader class of penalty functions.

The next lemma uses simple properties of length spaces (defined in Section 2) to bound the change resulting from deleting a member of S.

Lemma 5.1 *Let* $(\mathbb{U}, D)$ *be a length space. For finite set* $S \subset \mathbb{U}$*,* $p \in S$*, and* $x \in \mathrm{Vor}(p, S)$*, we have* $D(x, S \setminus \{p\}) \leq 3C(p, S)$*.*

Proof:
Suppose p' is the closest point of $S \setminus \{p\}$ to p. Then there is a point x' on the shortest path from p to p' such that $D(p, x') = D(p', x') = D(p, p')/2$. Moreover, $D(x', p'') \geq D(p, p')/2$ for $p'' \in S \setminus \{p, p'\}$, since otherwise $D(p, p'') \leq D(p, x') + D(x', p'') < D(p, p')$, by the triangle inequality. For any $x \in \mathrm{Vor}(p, S)$, again using the triangle inequality,

$$D(x, S \setminus \{p\}) \leq D(x, p') \leq D(x, p) + D(p, x') + D(x', p') \leq 3C(p, S),$$

and the lemma follows. ■

Say that metric measure space $Z = (\mathbb{U}, D, \mu)$ is *Voronoi regular* if there is threshold t_V such that for finite $S \subset \mathbb{U}$ with $D(\mathbb{U}, S) \leq t_V$, it holds that $\sum_{p \in S} \mu(\mathrm{Vor}(p, S)) = \mu(\mathbb{U})$.

Gruber[Gru04] shows that a nonzero threshold t_V exists for Riemannian manifolds, but the following, from Remark 16 of Chrusciel *et al.*[CFGH02], shows that t_V can be taken as infinite. It can be proven using the Lipschitz property of the function $D(x, S)$, Rademacher's result that Lipschitz functions are differentiable almost everywhere, and the nondifferentiability of $D(x, S)$ at perpendicular bisectors.

Lemma 5.2 *On a Riemannian manifold* $(\mathbb{U}, D)$ *and* $S \subset \mathbb{U}$*, the set of points in* $\mathbb{U}$ *that are equidistant from some two points in* S *has measure zero.*

Say that Z is *dimension regular* if there is a threshold t_D such that the balls $B(x, \epsilon)$ are measurable for all $x \in \mathbb{U}$ and $\epsilon < t_D$, and also

$$H(\epsilon) := \sup_{x \in \mathbb{U}} \mu(B(x, \epsilon))/\epsilon^d$$

and

$$L(\epsilon) := \inf_{x \in \mathbb{U}} \mu(B(x, \epsilon))/\epsilon^d$$

are in $\Theta(1)$ as $\epsilon \to 0$. (These values are related to the injectivity radius. This condition is stronger than the definition given by Cutler[Cut93] for dimension regularity, which requires only that $\log(\mu(B(x, \epsilon))/\epsilon^d)$ is in $o(\log(1/\epsilon))$. The condition is equivalent to the *positive density* condition of Gruber[Gru04], combined with the compactness of $\mathbb{U}$.)

The following lemma implies that we can assume that $D(\mathbb{U}, S_n)$ is small, for large enough n.

Lemma 5.3 *If metric measure space* $Z = (\mathbb{U}, D, \mu)$ *is dimension regular, then sets* $S_n \subset \mathbb{U}$ *of size* n *that minimize*

$$F(S) := \int D(x, S)^2 d\mu(x)$$

have $D(\mathbb{U}, S_n) = o(1)$ *as* $n \to \infty$*.*

Proof: Find an ϵ-net E of $\mathbb{U}$ for $\epsilon < t_D$, using for example the greedy algorithm. Since E is an ϵ-cover, $F(E) \leq \epsilon^2 \sum_{p\in E} \mu(B(p,\epsilon))$. Since E is an ϵ-packing, the balls $B(p,\epsilon/2)$ are disjoint. By the dimension regularity assumption, each $B(p,\epsilon/2)$ is measurable, and so by the disjointness and measurability, $\sum_{p\in E} \mu(B(p,\epsilon/2)) \leq \mu(\mathbb{U})$. Again by dimension regularity, $\mu(B(p,\epsilon/2)) = \Omega(\mu(B(p,\epsilon))$, and so $F(E) = \epsilon^2 O(\mu(\mathbb{U}))$. Since $F(S_{|E|}) \leq F(E)$,

$$F(S_n) = o(1)$$

as $n \to \infty$.

Moreover, for any n, if $q \in \mathbb{U}$ has $D(q, S_n) = D(\mathbb{U}, S_n)$, then the ball $\hat{B} := B(q,\hat{D})$, where $\hat{D} := \min\{t_D, D(\mathbb{U}, S_n)/3\}$, has $D(x, S_n) > 2D(\mathbb{U}, S_n)/3$ for all $x \in \hat{B}$, by the triangle inequality.

Thus $F(S_{|E|}) \geq (2D(\mathbb{U}, S_{|E|})/3)^2 \Omega(\hat{D}^d)$, and so $D(\mathbb{U}, S_n) = o(1)$ also as $n \to \infty$, as claimed. ∎

Theorem 5.4 *Suppose $Z = (\mathbb{U}, D, \mu)$ is a compact metric measure space, for length metric D and Borel measure μ. Suppose that Z is Voronoi regular and dimension regular. Then a set $S_n \subset \mathbb{U}$ of size n that minimizes*

$$F(S) := \int D(x,S)^2 d\mu(x)$$

is an $(\Omega(\beta(n)), O(\beta(n)))$-Delone set, where

$$\beta(n) := (\mu(\mathbb{U})/n)^{1/d},$$

as $n \to \infty$, and the constants depend only on d.

Proof: Suppose S is a set of size n with $D(\mathbb{U}, S)$ less than t_D and t_V; by the lemma just above, for large enough n this is a necessary condition for S to be an optimal set S_n.

To prove the covering and packing conditions, we will show that if they fail for such an S, then there is a point $p \in S$ and a point $q \in \mathbb{U}$ such that $F(S \setminus \{p\} \cup q)$ is smaller than $F(S)$, so that S cannot be optimal. We will pick a point p in this pivoting scheme to show that S is an $O(\beta(n))$-covering, and then pick a different p to show that S is an $\Omega(\beta(n))$-packing.

The point q will be one realizing $D(\mathbb{U}, S')$, where $S' := S \setminus \{p\}$, so

$$\hat{D} := D(q, S') = D(\mathbb{U}, S') \geq D(\mathbb{U}, S).$$

That is, S', and so S, are $\hat{D}$-coverings. Points in the ball $B(q, \hat{D}/3)$ have distance to q no more than $\hat{D}/3$, by definition, and at least $2\hat{D}/3$ from any point in S', by the triangle inequality. Therefore any point $x \in B(q, \hat{D}/3)$ has

$$D(x, S')^2 - D(x, S' \cup \{q\})^2 \geq (2\hat{D}/3)^2 - (\hat{D}/3)^2 = \hat{D}^2/3,$$

so

$$F(S') - F(S' \cup \{q\}) \geq \mu(B(q, \hat{D}/3))\hat{D}^2/3. \tag{4}$$

To show that S_n is a $O(\beta(n))$-covering, pick $p \in S$ that has the smallest Voronoi region in measure. Voronoi regularity implies that p has $\mu(\mathrm{Vor}(p, S)) \leq \mu(\mathbb{U})/n$. By Lemma 5.1, each $x \in \mathrm{Vor}(p, S)$ has $D(x, S') \leq 3C(p, S) \leq 3\hat{D}$. Thus deleting p increases F by at most $(3\hat{D})^2\mu(\mathbb{U})/n$, that is,

$$F(S) - F(S') \geq -(3\hat{D})^2\mu(\mathbb{U})/n.$$

By this fact, dimension regularity, and (4), we have

$$\begin{aligned} &F(S) - F(S \setminus \{p\} \cup \{q\}) \\ &\quad= F(S) - F(S') + F(S') - F(S' \cup \{q\}) \\ &\quad\geq -(3\hat{D})^2\mu(\mathbb{U})/n + (\hat{D}^2/3)\mu(B(q, \hat{D}/3)) \\ &\quad\geq \hat{D}^2[\Omega((\hat{D}/3)^d)/3 - 9\mu(\mathbb{U})/n], \end{aligned}$$

as $\hat{D} \to 0$. This expression is greater than zero when

$$\Omega((\hat{D}/3)^d) > 27\mu(\mathbb{U})/n.$$

Thus if $\hat{D} = D(\mathbb{U}, S')$ is not in $O(\mu(\mathbb{U})/n)^{1/d}$, S cannot be optimal for its size. Since S is an $\hat{D}$-covering, S_n must be an $O(\beta(n))$-covering as claimed.

To show that S_n must be an $\Omega(\beta(n))$-packing, assume as before that $D(\mathbb{U}, S)$ is less than t_D and t_V, and also S is an $O(\beta(n))$-covering; these are all necessary for optimality.

Now pick the closest pair of points $p, p' \in S$, at distance $\tilde{D} := D(p, p')$. We will show that if $\tilde{D}$ is too small, then $F(S \setminus \{p\} \cup q)$ is smaller than $F(S)$, so that S is not optimal, where q is chosen as before. Note that by an argument as in Lemma 5.1,

$$\tilde{D} \leq 2C(p, S) \leq 2D(\mathbb{U}, S) \leq 2D(\mathbb{U}, S \setminus \{p\}) =: \hat{D}. \tag{5}$$

For $x \in \mathrm{Vor}(p, S)$ we have $D(p', x) \leq \tilde{D} + D(p, x)$, so that

$$\begin{aligned} D(x, S')^2 - D(x, S)^2 &\leq (\tilde{D} + D(p, x))^2 - D(p, x)^2 \\ &= \tilde{D}(\tilde{D} + 2D(p, x)) \\ &\leq \tilde{D}(\tilde{D} + 2\hat{D}). \end{aligned} \tag{6}$$

Using (4), (5), (6), and dimension regularity,

$$\begin{aligned} F(S) - F(S \setminus \{p\} \cup \{q\}) &\geq (\hat{D}^2/3)\Omega((\hat{D}/3)^d) - \tilde{D}(\tilde{D} + 2\hat{D})O(\tilde{D}^d), \\ &\geq [\Omega(\hat{D}^2) - O(\tilde{D}(\tilde{D} + 2\hat{D}))]\hat{D}^d, \end{aligned}$$

as $\hat{D} \to 0$. This is greater than zero when $\tilde{D}$ is less than a constant factor times $\hat{D}$.

It holds that $\hat{D} = \Omega(\beta(n))$, which can be shown as follows. Since $\hat{D} \leq t_V$, there is some $p' \in S'$ with $\mu(\mathrm{Vor}(p', S')) \geq \mu(\mathbb{U})/(n-1)$. Letting $C_{p'} := C(p', S)$, we have $\mu(\mathrm{Vor}(p', S')) = O(C_{p'}^d)$, and so

$$\hat{D} \geq C_{p'} = \Omega((\mu(\mathbb{U})/(n-1))^{1/d}) = \Omega(\beta(n)).$$

So for optimal S_n, $\tilde{D}$ must at least a constant factor times $\hat{D}$, which is $\Omega(\beta(n))$. That is, the distance $\tilde{D}$ between the closest pair of points in S must be $\Omega(\beta(n))$, which means that S_n is an $\Omega(\beta(n))$-packing. ∎

This completes a version of the argument by Gruber. The following corollary implies that ϵ-nets give constant-factor approximations.

Corollary 5.5 *Under the conditions of Theorem 5.4, a greedy ϵ-net E of size n has $F(E) = O(F(S_n))$ as $n \to \infty$, where $F(S_n)$ is optimal for a set of size n. The constant factor is at most exponential in the manifold dimension d.*

Proof: From Theorem 5.4, S_n is a β-packing and a β_h-cover, for some β and $\beta_h = O(\beta)$. So

$$F(S_n) \geq n(\beta/2)^2 \Omega(\beta^d) = n\Omega(\beta^{d+2}),$$

while a greedy ϵ-net E of size n has

$$F(E) \leq n(\epsilon)^2 O(\epsilon^d) = nO(\epsilon^{d+2}).$$

The greedy construction has the property that the associated ϵ_n for n points is no more than twice the optimal cover radius for n points. Therefore $\epsilon_n \leq \beta_h = O(\beta)$, and the greedy ϵ-net has $F(E) = O(F(S_n))$ as $n \to \infty$, as claimed. ∎

The following characterization of optimal sets S_n may also be interest.

Theorem 5.6 *Under the conditions of Theorem 5.4, for an optimal set S_n and for any two points $p, \hat{p} \in S_n$, we have $C(\hat{p}, S_n) \leq \beta_d C(p, S_n)$, where*

$$\beta_d \leq 6(H(C(p, S_n))/L(C(\hat{p}, S_n)/3))^{d/(d+2)}.$$

That is, for a family of manifolds with $H(\epsilon)$ and $L(\epsilon)$ only exponentially dependent on d, the circumradius ratio is in not increasing in d.

Proof: We will show that if the condition does not hold, then $F(S \setminus \{p\} \cup \{q\})$ is smaller than $F(S)$, where $q \in \mathbb{U}$ realizes $C(\hat{p}, S \setminus \{p\}) \geq C(\hat{p}, S) =: \hat{D}$. As in the proof of the theorem above, the reduction in F due to adding q to S is at least $(\hat{D}^2/3)\Omega((\hat{D}/3)^d)$. Letting $\tilde{C} := C(p, S)$, $S' := S \setminus \{p\}$, and p' be closest in S' to p, for $x \in \mathrm{Vor}(p, S)$, the increase in distance is, using $D(p, p') \leq 2\tilde{C}$,

$$\begin{aligned} D(x, S')^2 - D(x, S)^2 &\leq D(p', x)^2 - D(p, x)^2 \\ &\leq (D(p, p') + D(p, x))^2 - D(p, x)^2 \\ &= D(p, p')(D(p, p') + 2D(p, x)) \leq 2\tilde{C}(2\tilde{C} + 2\tilde{C}) \leq 8\tilde{C}^2. \end{aligned}$$

Since also $\mu(\mathrm{Vor}(p, S)) \leq L(\tilde{C}) = O(\tilde{C}^d)$, we have

$$F(S) - F(S \setminus \{p\} \cup \{q\}) \geq (\hat{D}^2/3)(\hat{D}/3)^d L(\hat{D}/3) - 8\tilde{C}^2(\tilde{C})^d H(\tilde{C}),$$

so we must have $(\hat{D}/\tilde{C})^{d+2} \leq 3^{d+1} 8H(\tilde{C})/L(\hat{D}/3)$ for optimality, or $\hat{D} \leq \tilde{C} 6(H(\tilde{C})/L(\hat{D}/3))^{1/(d+2)}$. ∎

6 The curvature distance and Dudley's construction

As discussed by Pottmann *et al.* [PSH$^+$04] and mentioned in §1 and §3, the *regularized isophotic distance* has metric tensor $q_{\mathrm{I+III}} := q_{\mathrm{I}} + q_{\mathrm{III}}$; that is, it combines the metric tensor q_{I} for arc length and q_{III} for arc length in the Gauss map image. As discussed in §3, at a point p the matrix of $q_{\mathrm{I+III}}$ is $I + H_p^2$ in the coordinate system is oriented to p. We will simply write H for H_p below.

The analogous matrix for the root-curvature metric q_{II} is H. For any d-vector v, its measure by the isophotic metric tensor at p is $\sqrt{v^T v + v^T H^2 v}$, while for the root-curvature metric, it is $\sqrt{v^T H v}$. Using the Cauchy-Schwartz inequality and the arithmetic-geometric mean inequality, we have

$$\begin{aligned} q_{\mathrm{II}}(v;p) = v^T H v \le \|v\| \|Hv\| &\le (\|v\|^2 + \|Hv\|^2)/2 \\ &= (v^T v + v^T H^2 v)/2 = q_{\mathrm{I+III}}(v;p). \end{aligned}$$

Thus the isophotic distance between two points is always greater than the root-curvature distance, and so an isophotic ϵ-cover is always an ϵ-cover for the root-curvature distance.

By Lemma 3.1, the size of an ϵ-net for $D_{\mathrm{I+III}}$ is proportional to $1/\epsilon^d$ times

$$\begin{aligned} \int_M \sqrt{\det(I + H_x^2)}dx &= \int_M \sqrt{\prod_i (1 + k_i(x)^2)}dx \\ &\le \int_M \prod_i (1 + |k_i(x)|)dx, \end{aligned}$$

where $k_i(x)$ is the i'th eigenvalue of H. For a 2-manifold, this is no more than

$$\int_M dx + \int_M |k_1(x)| + |k_2(x)| dx + \int_M |k_1(x) k_2(x)| dx.$$

These terms are the surface area, the total mean curvature of the convexified metric, and the total Gaussian curvature of the convexified metric, respectively. The last term is more commonly known as the total absolute curvature. The total mean curvature of the convexified metric is within a constant factor of the "root mean square," or RMS, curvature, with an integrand of $\sqrt{k_1(x)^2 + k_2(x)^2}$. This bound can be strictly larger than for root-curvature distance.

As a concrete example, consider a circular cylinder of radius r, where the surface area of the "wrapped" part (not one of the capping disks) is 1. While the Gaussian curvature at a "wrapped" point is zero, the mean and RMS curvatures are $1/r$, and so the total RMS curvature is at least $1/r$. As $r \to 0$, this term in the bound dominates the surface area term. Thus the size of an ϵ-net in the $D_{\mathrm{I+III}}$ metric is proportional to $(1/r)/\epsilon^2$. In contrast, the size of an ϵ-net on the cylinder in the D_{II} metric is to $O(1/\epsilon^2)$, not increasing in $1/r$, where the cost comes from the capping disks.[3] Since these two nets yield triangulations that have Hausdorff error $O(\epsilon^2)$, the net for D_{II} is more economical for given error.

[3] Neglecting the δ-perturbation needed for positive-definiteness; on the other hand, a size

7 Metric Versions of Local Feature Size

While previous sections considered simplicial meshes for approximating a manifold, here the focus is on triangulations of planar regions, where a triangle is used for piecewise-linear interpolation of a function. A common condition on such meshes is that the triangles be as close to equilateral as possible. This implies that the meshes are *graded*, so that neighboring triangles are not too different in size.

Given a planar straight-line graph (PSLG) $G = (V, E)$, the *local feature size* $D(x, G)$ at point x is the smallest radius r so that the ball $B(x, r)$ meets an edge e_1 of G, and also meets another edge e_2 of G that does not meet e_1. Here an edge includes its endpoints, so edges e_1 and e_2 cannot share endpoints.

The motivation for this definition is that the local feature size at vertex v of a graded mesh is within a constant factor of the length of an edge incident to v. We could imagine the local feature size as providing a scaling to the distance between two points x and y: $D_E(x, y)/D(x, G)$ is a measure of the proximity of x and y, relative to G. This is not symmetrical in x and y; also, it doesn't really make any sense if x and y are far from each other. However, it is not hard to create a distance measure normalized by the local feature size that is symmetrical, and even obeys the triangle equality; that is, a distance measure that is a metric.

As shown below, the following two distance measures are metrics:

$$D_1(x, y) := \min\{1, \frac{2D(x, y)}{D(x, y) + D(x, G) + D(y, G)}\},$$

and

$$D_2(x, y) := \frac{2D(x, y)}{D(x, y) + \inf_{e_1, e_2 \in E, e_1 \cap e_2 = \{\}} D^t(x, y, e_1) + D^t(x, y, e_2)},$$

that is, e_1 and e_2 are edges of G that do not meet, and

$$D^t(x, y, e) := \frac{1}{2} \inf_{a \in e} D_E(x, a) + D_E(a, y),$$

the shortest path between x to y via e.

Both of these distance measures evidently have a maximum value of 1, and the only metric property that they do not obviously satisfy is the triangle inequality. It will also be shown that when the distances are small, say less than $1/5$, that they are reasonably well approximated by $D(x, y)/D(x, G)$.

These constructions begin from the *biotope transform*, also called the *Steinhaus transform*:

more like $O(\sqrt{\alpha r}/\epsilon^2)$ as $\epsilon \to 0$ would be obtained if the cylinder were approximated by the convex hull of two torii of major radius r and minor radius α. The Gaussian curvature on a torus in the region of interest is no more than $1/\alpha r$, and the area is αr, so the total root curvature contributed by the torii is proportional to $\sqrt{\alpha r}$.

Lemma 7.1 *[DLD97, Prop. 9.2.1] If* $(\mathbb{U}, D)$ *is a metric space,* $a \in \mathbb{U}$*, and the distance measure* $D_{\{a\}}$ *is defined by*

$$D_{\{a\}}(x,y) := \frac{D(x,y)}{D(x,y) + D(x,a) + D(y,a)},$$

then $(\mathbb{U}, D_{\{a\}})$ *is also a metric space.*

Note that if $(\mathbb{U}, D)$ is a metric space, then so is $(\{a,x,y,z\}, D)$, where $\{a,x,y,z\} \subset \mathbb{U}$.

To show that D_1 satisfies the triangle inequality, we use the well-known property that $D(x,G)$ is 1-Lipschitz, so that $D(y,G) \leq D(x,y) + D(x,G)$ for all x and y. The claim then follows from the following more general one.

Lemma 7.2 *Let* $(\mathbb{U}, D)$ *be a metric space, and* $F : \mathbb{U} \to \Re^+$ *be a 1-Lipschitz function, so* $|F(x) - F(y)| \leq D(x,y)$ *for all* $x, y \in \mathbb{U}$*. Then*

$$D_F(x,y) := \min\{1, \frac{2D(x,y)}{D(x,y) + F(x) + F(y)}\}$$

is a metric on $\mathbb{U}$*.*

Note that the construction can be generalized to a β-Lipschitz function F that has an infinum, since $(F(x) - \inf_{y \in \mathbb{U}} F(y))/\beta$ is a nonnegative 1-Lipschitz function.

Proof: All properties except the triangle inequality are trivial to verify. For the latter, note that $D_F(x,z) \leq 1$ by construction, so that if $D_F(x,y) = 1$ or $D_F(y,z) = 1$, then immediately $D_F(x,z) \leq D_F(x,y) + D_F(y,z)$.

So suppose $D_F(x,y) < 1$ and $D_F(y,z) < 1$. From the definition, this implies

$$F(x) + F(y) > D(x,y) \text{ and } F(y) + F(z) > D_F(y,z) \tag{7}$$

If also $D_F(x,z) < 1$, then

$$F(x) + F(z) > D(x,z). \tag{8}$$

Define the distance to some adjoined object a by $D(w,a) := F(w)$ for $w \in \{x,y,z\}$, and of course $D(a,a) := 0$. Then (7) and (8), together with the 1-Lipshitz property, imply that $(\{a,x,y,z\}, D)$ is a metric space. (If $D(w,a) = F(w) = 0$ for some $w \in \{x,y,z\}$, then $D(w,w') = 1$ for any w', so $D(w,a) \neq 0$ here for $w \neq a$.) Therefore $(\{a,x,y,z\}, D_{\{a\}})$ is the biotope transform of $(\{a,x,y,z\}, D)$, and the triangle inequality holds for $D_{\{a\}}$ at x, y, and z. Since $D_{\{a\}}$ agrees with D_F at x, y, and z, the triangle inequality holds also for D_F at x, y, and z, when $D_F(x,z) < 1$.

Now suppose $D_F(x,z) = 1$, but still $D_F(x,y) < 1$ and $D_F(y,z) < 1$. Then $F(x) + F(z) \leq D(x,z)$. This condition and the 1-Lipschitz property imply

$$\begin{aligned} D(x,y) + F(x) + F(y) &\leq D(x,y) + F(x) + (F(z) + D(y,z)) \\ &\leq D(x,y) + F(x) + (D(x,z) - F(x) + D(y,z)) \\ &= D(x,y) + D(x,z) + D(y,z), \end{aligned}$$

and similarly

$$D(y,z)+F(y)+F(z) \leq D(y,z)+D(z,x)+D(y,x) = D(x,y)+D(x,z)+D(y,z).$$

Thus

$$\begin{aligned}
D_F(x,y) + D_F(y,z) \\
&= \frac{2D(x,y)}{D(x,y)+F(x)+F(y)} + \frac{2D(y,z)}{D(y,z)+F(y)+F(z)} \\
&\geq \frac{2D(x,y)}{D(x,y)+D(x,z)+D(y,z)} + \frac{2D(y,z)}{D(x,y)+D(x,z)+D(y,z)} \\
&= \frac{2(D(x,y)+D(y,z))}{D(x,y)+D(x,z)+D(y,z)} \\
&\geq 1 = D_F(x,z),
\end{aligned}$$

since $D(x,z) \leq D(x,y) + D(y,z)$.

Thus whether $D_F(x,z) < 1$ or not, the triangle inequality holds for D_F. ■

The proof of the triangle inequality for D_2 proceeds from the following generalization of the biotope transform.

Lemma 7.3 *Suppose* $(\mathbb{U}, D)$ *is a metric space,* $Q \subset \mathbb{U}$ *is closed, and*

$$D_Q(x,y) := \frac{D(x,y)}{D(x,y) + \inf_{a \in Q} D(x,a) + D(y,a)}.$$

Then $(\mathbb{U}, D_Q)$ *is also a metric.*

Proof: The only property that is not immediate is the triangle inequality. For given $x,y,z \in \mathbb{U}$, suppose $\hat{a} \in Q$ yields $\inf_{a \in Q} D(x,a) + D(z,a)$, so that $D_Q(x,z) = D_{\{\hat{a}\}}(x,z)$. Then by Lemma 7.1,

$$D_Q(x,z) = D_{\{\hat{a}\}}(x,z) \leq D_{\{\hat{a}\}}(x,y) + D_{\{\hat{a}\}}(y,z) \leq D_Q(x,y) + D_Q(y,z),$$

where the last inequality follows from the definition of D_Q. So D_Q obeys the triangle inequality. ■

Note that

$$\min\{1, \frac{D(x,y)}{D(x,y) + \inf_{a \in Q} D(x,a) + \inf_{a \in Q} D(y,a)}\}$$

is also a metric, since $\inf_{a \in Q} D(x,a)$ is 1-Lipschitz, but is not quite the same as D_Q.

Example: relation to hyperbolic metrics. Consider the points in the unit disk $U := \{(x,y) \mid x^2 + y^2 \leq 1\}$, let D be the Euclidean distance, and let C be the unit circle $\operatorname{bd} U$. Then

$$D_C(p_1,p_2) = \frac{D(p_1,p_2)}{1 - \|p_1\|}(1 + O(D(p_1,p_2))$$

as $D(p_1, p_2) \to 0$. The metric tensor of the Poincaré disk model of hyperbolic geometry at point z is

$$ds^2 = \frac{dx^2 + dy^2}{1 - z^2}.$$

It seems that D_C behaves roughly like the hyperbolic tangent of the Poincaré metric. It remains to be seen if (U, D_C) is a hyperbolic metric space.

Theorem 7.4 *Given a PSLG $G = (V, E)$, the distance measures $D_1(x, y)$ and $D_2(x, y)$, for $x, y \in \Re^2$, are metrics.*

Proof: Again, since $D(x, G)$ is 1-Lipschitz, Lemma 7.2 immediately implies that $D_1(x, y) = D_{D(\cdot,G)}(x, y)$ satisfies the triangle inequality, and the remaining conditions for a metric are trivial to show.

The triangle inequality for D_2 follows from the generalization of the biotope transform in the lemma just above. Let

$$Q := \{(p_1, p_2) \mid p_1 \in e_1, p_2 \in e_2, e_i \in E, e_1 \cap e_2 = \{\}\}.$$

For $x, y \in (\Re^2)^2$, so $x = (x_1, x_2)$ and $y = (y_1, y_2)$ with $x_1, x_2, y_1, y_2 \in \Re^2$, let $D(x, y) := D_E(x_1, y_1) + D_E(x_2, y_2)$. Let $P := \{(x, x) \mid x \in \Re^2\}$, and let $\mathbb{U} = Q \cup P$. Then $(\mathbb{U}, D)$ is a metric space, and D_Q is a metric, by the previous lemma. Moreover, for $x, y \in P$,

$$\begin{aligned}
&D_Q(x, y) \\
&= \frac{2D(x, y)}{D(x, y) + \inf_{a \in Q} D(x, a) + D(y, a)} \\
&= \frac{4D_E(x_1, y_1)}{2D_E(x_1, y_1) + \inf_{\substack{a_1 \in e_1 \\ a_2 \in e_2}} D_E(x_1, a_1) + D_E(x_1, a_2) + D_E(y_1, a_1) + D_E(y_1, a_2)} \\
&= \frac{2D_E(x_1, y_1)}{D_E(x_1, y_1) + \inf_{e_1 \cap e_2 = \{\}} D^t(x_1, y_1, e_1) + D^t(x_1, y_1, e_2)} \\
&= D_2(x, y).
\end{aligned}$$

So D_1 and D_2 satisfy the triangle inequality, and are metrics. ■

Lemma 7.5 *Given a metric space $(\mathbb{U}, D)$, $A \subset \mathbb{U}$, and 1-Lipschitz function F on $\mathbb{U}$,*

$$D_F(x, y) = \frac{D(x, y)}{F(x)}((1 - D_F(x, y)(1 + \gamma)/2)),$$

where $0 \le \gamma \le 1$, and also

$$D_A(x, y) = \frac{2D(x, y)}{D_m}(1 - D_A(x, y)/2),$$

where $D_m = \inf_{a \in A} D(x, a) + D(y, a)$.

Proof: Dropping the E from D_E, we have

$$\frac{2D(x,y)}{D(x,y)+F(x)+F(y)} = D_1(x,y),$$

or

$$\begin{aligned} 2D(x,y) &= D_F(x,y)(D(x,y)+F(x)+F(y)) \\ &= D_F(x,y)(D(x,y)+F(x)+F(x)+\gamma D(x,y)), \end{aligned}$$

for some γ with $0 \leq \gamma \leq 1$, or

$$(2 - D_F(x,y)(1+\gamma))D(x,y) = 2D_F(x,y)F(x)$$

so

$$D_F(x,y) = \frac{D(x,y)}{F(x)}(1 - D_1(x,y)(1+\gamma)/2).$$

A similar argument gives the result for D_A.

■

8 Interpretation of Previous Results

By applying the greedy algorithm for ϵ-nets to the local feature metrics D_1 or D_2, then computing the (Euclidean) Delaunay triangulation of the ϵ-net, the result is a graded triangulation. By first picking an ϵ'-net for each edge of G individually, with ϵ' sufficiently smaller than ϵ, a point set can be constructed so that the triangulation conforms to G, that is, the edges of G are the union of edges of the triangulation. The resulting triangulation has triangles that are "nicely shaped".

Such an algorithm is very close to that of Chew[Che89], or of Ruppert[Rup95], but may give somewhat smoother results.

In surface reconstruction, sample points on a manifold are known, and an approximation to the unknown manifold is desired. Several papers on surface reconstruction, starting with that of Amenta and Bern[AB99], give sufficient conditions on the sample points so that a well-behaved reconstruction of the manifold is possible. These conditions use a "local feature size" that involves distance to the medial axis of the manifold. The conditions amount to the requirement that the sample points S are a Delone set with respect to the metric D_A which is the Euclidean distance, scaled by the distance to the medial axis. The conditions can require sample sets that are larger than those required for an approximating mesh. Of course, the requirements on the sample for surface reconstruction are more stringent, and the resulting mesh satisfies consistency conditions that the asymptotic approach here satisfies only for ϵ sufficiently small.

9 Acknowledgement.

I am grateful to Gopi Meenakshisundaram for pointing out an error in the example of Section 6.

References

[AB99] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete and Computational Geometry*, 22:481–504, 1999.

[ACK01] Nina Amenta, Sunghee Choi, and Ravi Kolluri. The power crust, unions of balls, and the medial axis transform. *Computational Geometry: Theory and Applications*, 19:127–153, 2001.

[AGK+01] Vijay Arya, Naveen Garg, Rohit Khandekar, Kamesh Munagala, and Vinayaka Pandit. Local search heuristic for k-median and facility location problems. In *STOC '01: Proc. Thirty-Third Annual ACM Symp. on Theory of Computing*, pages 21–29, 2001.

[AHPV05] P.K. Agarwal, S. Har-Peled, and K. Varadarajan. Geometric approximation via coresets. In E. Welzl, editor, *Current Trends in Combinatorial and Computational Geometry.* Cambridge University Press, 2005.

[CFGH02] P. T. Chruściel, J.H.G. Fu, G. Galloway, and R. E. Howard. On fine differentiability properties of horizons and applications to Riemannian geometry. *J. Geometry and Physics*, 41:1–12, 2002.

[Che89] L. Paul Chew. Guaranteed-quality triangular meshes. Technical Report TR-89-983, Cornell University, 1989.

[Cla94] K. L. Clarkson. An algorithm for approximate closest-point queries. In *SOCG '94: Proc. Tenth Annual ACM Symp. on Computational Geometry*, pages 160–164, 1994.

[CSAD04] David Cohen-Steiner, Pierre Alliez, and Mathieu Desbrun. Variational shape approximation. *ACM Trans. Graph.*, 23(3):905–914, 2004.

[CSXar] L. Chen, P. Sun, and J. Xu. Optimal anisotropic simplicial meshes for minimizing interpolation errors in l_p-norm. *Math. Comp.*, To Appear.

[Cut93] C. D. Cutler. A review of the theory and estimation of fractal dimension. In H. Tong, editor, *Dimension Estimation and Models.* World Scientific, 1993.

[DFG99] Qiang Du, Vance Faber, and Max Gunzburger. Centroidal Voronoi tessellations: Applications and algorithms. *SIAM Rev.*, 41(4):637–676, 1999.

[DLD97] M. Deza, M. Laurent, and M. M. Deza. *Geometry of Cuts and Metrics*. Springer, 1997.

[Dud74] R. M. Dudley. Metric entropy of some classes of sets with differentiable boundaries. *J. Approximation Theory*, 10:227–236, 1974.

[ELPZ97] Y. Eldar, M. Lindenbaum, M. Porat, and Y. Zeevi. The farthest point strategy for progressive image sampling, 1997.

[Eri] Jeff Erickson. Personal Communication.

[Gon85] T. Gonzalez. Clustering to minimize the maximum intercluster distance. *Theoret. Comput. Sci.*, 38:293–306, 1985.

[Gru93] P. M. Gruber. Asymptotic estimates for best and stepwise approximation of convex bodies I. *Forum Math.*, 5:281–297, 1993.

[Gru04] P. M. Gruber. Optimum quantization and its applications. *Adv. Math.*, 186:456–497, 2004.

[HG99] Paul S. Heckbert and Michael Garland. Optimal triangulation and quadric-based surface simplification. *Comput. Geom. Theory Appl.*, 14(1-3):49–65, 1999.

[LL00] Greg Leibon and David Letscher. Delaunay triangulations and Voronoi diagrams for Riemannian manifolds. In *SCG '00: Proceedings of the Sixteenth Annual Symposium on Computational geometry*, pages 341–349, New York, NY, USA, 2000. ACM Press.

[Lud98] M. Ludwig. Asymptotic approximation of convex curves; the Hausdorff metric case. *Arch. Math.*, 70:331–336, 1998.

[Nad86] E. Nadler. Piecewise linear best ℓ_2 approximation on triangulations. In C. K. Chui, L. L. Schumaker, and J. D. Ward, editors, *Approximation Theory V*, pages 499–502. Academic Press, 1986.

[PC05] Gabriel Peyré and Laurent D. Cohen. Geodesic computations for fast and accurate surface remeshing and parameterization. *Progress in Nonlinear Differential Equations and Their Applications*, 63:157–171, 2005.

[PKH+00] H. Pottmann, R. Krasauskas, B. Hamann, K. Joy, and W. Seibold. On piecewise linear approximation of quadratic functions. *J. Geom. Graphics*, 4:31–53, 2000.

[PSH+04] H. Pottmann, T. Steiner, M. Hofer, C. Haider, and A. Hanbury. The isophotic metric and its application to feature sensitive morphology on surfaces. In T. Pajdla and J. Matas, editors, *Computer Vision — ECCV 2004, Part IV*, volume 3024 of *Lecture Notes in Computer Science*, pages 560–572. Springer, 2004.

[Rup95] Jim Ruppert. A Delaunay refinement algorithm for quality 2-dimensional mesh generation. In *SODA '93: Proc. Fourth Annual ACM-SIAM Symp. on Discrete Algorithms*, pages 548–585, Orlando, FL, USA, 1995. Academic Press, Inc.

[She02] J. R. Shewchuk. What is a good linear finite element? Interpolation, conditioning, anisotropy, and quality measures. `http://www.cs.berkeley.edu/~jrs/papers/elemj.pdf`, 2002.