Manifold Reconstruction from Point Samples

Siu-Wing Cheng[*] Tamal K. Dey[†] Edgar A. Ramos[‡]

Abstract

We present an algorithm to "reconstruct" a smooth k-dimensional manifold $\mathcal{M}$ embedded in an Euclidean space $\mathbb{R}^d$ from a "sufficiently dense" point sample from the manifold. The algorithm outputs a simplicial manifold that is homeomorphic and geometrically close to $\mathcal{M}$. The running time is $O(n \log n)$ where n is the number of points in the sample (the multiplicative constant depends exponentially on the dimension though).

1 Introduction

There are a number of applications in science and engineering that deal with data points lying on a manifold embedded[1] in an Euclidean space. Data collected for scientific analysis through natural phenomena or simulations lie on such manifolds. This has led to the problem of *manifold learning* that, ideally, seeks to approximate a manifold embedded in an Euclidean space from a point sample [6, 17, 19]. Often, in practice, the data points are approximated with an appropriate linear flat. Principal component analysis (PCA) and the multidimensional scaling (MDS) are two prevalent methods used for this probelm [16, 18]. Although these are useful techniques, they are not appropriate to approximate points that come from a non-linear class such as smooth manifolds.

The cases for two and three dimensions where the manifold has co-dimension one have recently been solved [1, 2, 3, 7, 10] starting with the work of Amenta, Bern and Eppstein [2] in two dimensions and Amenta and Bern [1] in three dimensions. Dey, Giesen, Goswami and Zhao [11] gave an algorithm that detects the dimension of a manifold in any Euclidean space from its point sample. This algorithm also can reconstruct a curve or a surface in three dimensions thus solving the general problem in three dimensions. Giesen and Wagner [15] gave an improved algorithm in terms of time and space complexities. However, these algorithms cannot reconstruct the manifolds in higher dimensions in the sense of producing an approximation manifold that is topologically equivalent and geometrically close to the sampled one.

In this paper we present an algorithm that solves the general manifold reconstruction problem. Specifically, given a sufficiently dense point sample S of a smooth k-dimensional manifold $\mathcal{M} \subset \mathbb{R}^d$, the algorithm can produce a triangulation T interpolating S so that $\mathcal{M}$ and the underlying space of T, denoted $|T|$, are homeomorphic. Furthermore, the Hausdorff distance between $\mathcal{M}$ and $|T|$, and the (appropriate) distance between their respective normal spaces are provably small. The algorithm runs in time $O(n \log n)$ (the hidden constant depends exponentially on the dimension of the space). The algorithm builds on the *cocone* algorithm for surface reconstruction by Amenta, Choi, Dey and Leekha [3] and the sliver exudation technique to remove certain type of flat simplices, called *slivers*, from a simplicial mesh by Cheng, Dey, Edelsbrunner, Facello and Teng [8]. Intuitively, the Delaunay complex restricted to $\mathcal{M}$ (see definitions in next section) is a good candidate for a reconstruction. However, as $\mathcal{M}$ is not known, it cannot be directly computed: the difficulty lies in that Voronoi cells dual to $(k+1)$-dimensional simplices may lie near $\mathcal{M}$ and then make the extraction of $\mathrm{Del}_{\mathcal{M}}(S)$ ambiguous and difficult. The idea borrowed from [8] is that an assignment of weigths $\widehat{S}$ to the points, as a way of perturbation, moves any such Voronoi cell away from $\mathcal{M}$ and then extracting $\mathrm{Del}_{\mathcal{M}}(\widehat{S})$ from the Delaunay complex becomes essentially trivial.

The algorithm is not practical, mainly because it requires a very dense and noise-free sample, but also because it makes use of (weighted) Delaunay triangulation in higher dimensions (but note that in the final $O(n \log n)$ time algorithm, not all of the triangulation need to be computed). It is, however, the first such algorithm solving the general embedded manifold reconstruction problem. In proving the correctness

[*]Research was supported by the Research Grant Council, Hong Kong, China, projects HKUST6190/02E and HKUST 6169/03E. Department of Computer Science, The Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong.

[†]Research is supported by NSF CARGO grant DMS-0310642. Department of Computer Science and Engineering, The Ohio State University, Columbus, Ohio, USA.

[‡]Department of Computer Science, University of Illinois at Urbana Champaign, Urbana, Illinois, USA.

[1]A more general situation, that we do not consider, is *immersion*, in which the manifold is allowed to selfintersect.

of the algorithm, we extend several ideas in sliver exudation and surface reconstruction to higher dimensions, which are interesting on their own. The main theorem to be proved is (precise definitions follow later):

THEOREM 1.1. *There exist $\varepsilon > 0$ such that if S is a (ε, δ)-sampling of a compact k-manifold $\mathcal{M} \subset \mathbb{R}^d$ with a positive local feature size, then there is a weight assignment to S such that $\mathrm{Del}_{\mathcal{M}} \widehat{S}$ is a "faithful" reconstruction of $\mathcal{M}$ and can be "efficiently" computed.*

Faithful means that it has the right topology and it is a close geometric approximation (both in normal and distance) of $\mathcal{M}$. The ε obtained in the analysis is exponentially small in the dimension. Efficiently means time $O(n \log n)$, but the hidden constant is very large.

Contents. Section 2 contains definitions and some preliminary facts. The basic algorithm is described in Section 3. Sections 4-8 develop geometric facts needed to extend the method of sliver exudation. Sections 10-12 develop further geometric facts needed to verify the correctness of the output produced by the algorithm. Section 13 describes how the algorithm is modified to achieve the running time $O(n \log n)$, and how it can be extended to other less restrictive sampling conditions. Because of space limitation, several proofs have been omitted. They can be found in an extended version of this paper.

2 Definitions and Preliminaries

Weighted points. A weighted point $\widehat{p} \in \mathbb{R}^d \times \mathbb{R}^+$ is a d-ball with the center $p \in \mathbb{R}^d$ and weight (radius) $P \in \mathbb{R}^+$; we write $\widehat{p} = (p, P)$. *Unweighted* points correspond to points with zero radius. The *weighted distance* of $\widehat{x}$ from $\widehat{p}$ is $\pi_{\widehat{p}}(\widehat{x}) = \|x-p\|^2 - X^2 - P^2$. If $\pi_{\widehat{p}}(\widehat{x}) = 0$, we say that $\widehat{p}$ and $\widehat{x}$ are *orthogonal*. If $\pi_{\widehat{p}}(\widehat{x})$ is greater (smaller) than 0, we say that $\widehat{x}$ is *further (closer) than orthogonal* from $\widehat{p}$. The *bisector plane* of $\widehat{p}$ and $\widehat{q}$ is the locus of (unweighted) points x at equal weighted distances from $\widehat{p}$ and $\widehat{q}$, that is, $\pi_{\widehat{p}}(x) = \pi_{\widehat{q}}(x)$. Equivalently, x is in the bisector if there is a weight X such that $\widehat{x} = (x, X)$ is orthogonal to $\widehat{p}$ and $\widehat{q}$. We use $\widehat{S}$ to denote the set of points S with certain assignment of weights. We say that $\widehat{S}$ has *weight property* $[\omega]$ if for each point $p \in S$, $P \leq \omega \cdot N(p)$, where $N(p)$ is the distance to the closest point in S different from p.

Weighted Voronoi and Delaunay. We assume points in *general position*. For a set of weighted points $\widehat{S}$, the weighted Voronoi cell $V_{\widehat{s}}$ of $\widehat{s} \in \widehat{S}$ is $V_{\widehat{s}} = \{x \in \mathbb{R}^d \mid \pi_{\widehat{s}}(x) \leq \pi_{\widehat{r}}(x) \text{ for all } \widehat{r} \in \widehat{S}\}$. The weighted Voronoi cells and their k-dimensional faces, $0 \leq k \leq d$ form a polyhedral complex, denoted Vor $\hat{S}$, that decomposes $\mathbb{R}^d$. The $(d-j)$-dimensional faces are the nonempty cells $V_{\widehat{T}} = \bigcap_{\widehat{s} \in \widehat{T}} V_{\widehat{s}}$, where $\widehat{T} \subseteq \widehat{S}$ with $|\widehat{T}| = j$ (so both $V_{\widehat{s}}$ and $V_{\{\widehat{s}\}}$ denote the Voronoi cell of $\widehat{s}$). The weighted Delaunay triangulation $\mathrm{Del}\, \widehat{S}$ of $\widehat{S}$ is the *dual* to the weighted Voronoi diagram of $\widehat{S}$. That is, a simplex τ is in $\mathrm{Del}\, \widehat{S}$ iff $V_{\mathrm{vert}(\tau)} \neq \emptyset$, where $\mathrm{vert}(\tau)$ is the set of vertices of τ. Alternatively, a simplex τ is in $\mathrm{Del}\, \widehat{S}$ if there is a sphere (weighted points), called *orthosphere*, orthogonal to each weighted vertex and further than orthogonal from all other weighted points in $\widehat{S}$. For simplicity of notation, we write V_τ instead of $V_{\mathrm{vert}(\tau)}$, when τ is known. Also for simplicity, other variations in the notation appear later; for example V_{qr} is $V_{\{q,r\}}$. The meaning will be clear in the context. Vor S and $\mathrm{Del}\, S$ refer to the unweighted Voronoi and Delaunay complexes.

Sampled Manifold. We assume that the sampled k-dimensional manifold $\mathcal{M} \subset \mathbb{R}^d$ is compact, smooth and has no boundary. For any point $x \in \mathcal{M}$, $\mathcal{T}_x$ and $\mathcal{N}_x$ denote the tangent space and the normal space at x, respectively. The medial axis of $\mathcal{M}$ is the closure of the centers of the maximal balls that meet $\mathcal{M}$ only tangentially at two or more points. These balls are called *medial balls*. The *local feature size* $f(x)$ at a point $x \in \mathcal{M}$ is the distance of x to the medial axis of $\mathcal{M}$. We require that $f(x) > 0$ for all $x \in \mathcal{M}$. The local feature size is 1-Lipschitz, that is, $f(x) \leq f(y) + \|x - y\|$ for any two points x, y in $\mathcal{M}$.

(ε,δ)-Sampling. A set S of points on $\mathcal{M}$ is a (ε, δ)-sampling [11] if (i) for each point $x \in \mathcal{M}$, there is a sample $p \in S$ such that $\|p - x\| \leq \varepsilon f(x)$, and (ii) for any $p, q \in S$, $\|p-q\| \geq \delta f(p)$. We assume that ε and δ is within a constant factor. That is, ε/δ is a constant. An alternative condition is (ε, ℓ)-sampling in which we have (ii') for any $p \in S$, the number of q with $\|p - q\| \leq \varepsilon f(p)$ is at most ℓ, instead of (ii). This allows for some but limited locally dense sampling.

Restricted Voronoi and Delaunay. The restricted (weighted) Voronoi cell $V_{\mathcal{M},p}$ is given by $V_{\mathcal{M},p} = V_p \cap \mathcal{M}$. A simplex τ with the vertex set R is in the restricted (weighted) Delaunay triangulation $\mathrm{Del}_{\mathcal{M}} \widehat{S}$ if the restricted Voronoi cells of the vertices in R have a non-empty intersection. An important fact [12] that we use later is that if each restricted Voronoi cell is topologically a ball, then $\mathrm{Del}_{\mathcal{M}} \widehat{S}$ is homeomorphic to $\mathcal{M}$.

Cocone. If the sampling is dense, then for a sample point p, its neighbors in the sample relevant for creating a reconstruction can be found near $\mathcal{T}_p$. More precisely, within a cone of aperture θ_0 around $\mathcal{T}_p$ where $\theta_0 = O(\varepsilon)$. This can also be viewed as the complement of a cone – *cocone* – around $\mathcal{N}_p$

(and we follow [11] in refering to it as cocone). Algorithmically, the cocone is determined as follows [11]. First, a way to estimate $\mathcal{T}_p$ is needed. The *Voronoi subpolytopes* for a sample point $p \in \mathcal{M}$ are special subsets $V_p^i \subseteq V_p, i = 1, .., d$ of the Voronoi cell V_p, defined inductively as follows. For the base case let $V_p^d = V_p$. Inductively, assume that V_p^i is already defined. Let v_p^i be the farthest point in V_p^i from p. We call v_p^i the *pole* of V_p^i and the vector $\mathbf{v}_p^i = v_p^i - p$ its *pole vector*. If V_p^i is unbounded, v_p^i is taken at infinity, and the direction of $\mathbf{v}_p^i$ is taken as the average of all directions given by unbounded edges. The Voronoi subpolytope V_p^{i-1} is the intersection of V_p^i and the hyperplane $(x-p) \cdot \mathbf{v}_p^i = 0$. It is shown in [11] that $\mathrm{aff}(V_p^k)$ approximates $\mathcal{T}_p$ within an angle $O(\varepsilon)$. This carries over for a weighted Voronoi diagram given the weight property. (In the definitions here, V_p stands for the Voronoi cell of the weighted point p.) If k is not known in advance, then it can actually be determined through this construction, for appropriately uniform sampling [11]. Let $p \in S$ be a sample point from $\mathcal{M}$ where $\mathcal{M}$ has dimension k. The *cocone* of p, C_p, is defined as the set of all points $x \in V_p$ so that the segment connecting x and p makes an acute angle less than $\frac{\pi}{32}$ with V_p^k.A simplex τ in $\mathrm{Del}\, S$ is a C_p*-simplex,* if V_τ intersects C_p. In particular, if τ is an edge pq, we call pq a C_p*-edge* and q a *cocone neighbor* of p.

Simplex Shape. The *circumsphere* of a simplex τ is the smallest circumsphere of its vertices; its *circumcenter* and its *circumradius*, denoted R_τ, are the center (which lies in $\mathrm{aff}(\tau)$) and circumradius of the circumsphere. The length of the shortest edge of τ is denoted L_τ, The *circumradius-edge* ratio R_τ / L_τ is a measure of the quality of τ. Still, a simplex τ with $O(1)$ circumradius-edge ratio may be "very flat". This is captured by the concept of sliver simplices. For a $(j-1)$-simplex τ' and a point p not in $\mathrm{aff}(\tau')$, the *join* $p * \tau'$ of p and τ' is the convex hull of p and τ', a j-simplex. For a j-simplex τ and a vertex p of τ, let τ_p be the $(j-1)$-dimensional boundary simplex of τ such that $\tau = p * \tau_p$. The property of being sliver is quantified with a constant σ, to be chosen later. We define slivers by induction on the dimension: 0- and 1-simplices are not slivers; for $j \geq 2$, a j-simplex τ is a j*-sliver* if none of its boundary simplices is a sliver and for some vertex p of τ, $\mathrm{vol}(\tau) < \sigma L_\tau \mathrm{vol}(\tau_p)$ holds. Thus, if neither τ nor any of its boundary simplices is a sliver, then $\mathrm{vol}(\tau) \geq \sigma L_\tau \mathrm{vol}(\tau_p)$ for every vertex p of τ.

(ε, δ)**-Sampling Properties.** We henceforth assume that $\mathcal{M}$ is a smooth manifold with nonzero local feature size and that S is an (ε, δ)-sample of $\mathcal{M}$. Given a line segment pq, $\angle(pq, \mathcal{T}_p)$ denotes the angle between pq and $\mathcal{T}_p$. We will use the following result due to Giesen and Wagner [15].

LEMMA 1.

(i) The distance between $p \in S$ and its nearest neighbor in $S - \{p\}$ is at most $\frac{2\varepsilon}{1-\varepsilon} f(p)$.

(ii) For any points $p, q \in \mathcal{M}$ such that $\|p - q\| = t f(p)$ for some $0 < t < 1, \sin \angle(pq, \mathcal{T}_p) \leq t/2$.

(iii) Let p be a point in $\mathcal{M}$. Let q be a point in $\mathcal{T}_p$ such that $\|p - q\| \leq t f(p)$ for some $0 < t \leq 1/4$. Let q' be the point on $\mathcal{M}$ closest to q. Then $\|q - q'\| \leq 2t^2 f(p)$.

3 Algorithm

In this section, we present the basic reconstruction algorithm. Section 13 describes how to modify it to achieve running time $O(n \log n)$. The input is an (ε, δ)-sampling from the manifold $\mathcal{M}$ with ε sufficiently small (to be determined in the analysis). The algorithm proceeds as follows:

1. Construct $\mathrm{Vor}\, S$ and $\mathrm{Del}\, S$.
2. Determine the dimension k of $\mathcal{M}$.
3. "Pump up" the sample point weights to remove all j-slivers, $j = 3, \ldots, k+1$, from all point cocones.
4. Extract all cocone simplices as the resulting output.

Step 1 uses any standard algorithm, step 2 uses the algorithm in [11], and step 4 is essentially trivial (simply select the cocone simplices according to the definition). We elaborate now on step 3. This follows [8]. The algorithm iteratively assigns real weights to the sample points in an arbitrary order. Let p be the sample point being processed. As in sliver exudation [8], using the σ predicted by the theoretical result in section 8 would be extremely pessimistic. Instead, for each j-dimensional C_p-simplex τ, $3 \leq j \leq k+1$, we compute the interval $W(\tau)$ of weight for which τ appears in the triangulation and the sliver measure

$$\sigma(\tau) = \min_q \frac{\mathrm{vol}(\tau)}{L_\tau \mathrm{vol}(\tau_q)},$$

where the min is over all vertices of τ, We include all C_p-simplices in the weighted Delaunay triangulations generated by varying the weight of p and keeping all other weights fixed. The weight (radius) of p is varied within the range $[0, \omega N(p)]$. As the weight is increased from zero, all simplices in the weighted Delaunay triangulations incident to p are generated by repeated flip operations. For each simplex τ currently incident to p, we keep it in a priority queue indexed by the weight of p at which τ will be destroyed for the first time. Thus the minimum weight in the priority queue tells us the next event time. It is easy to extract the C_p-simplices from the entire set of simplices incident to p.

After we generate all C_p-simplices in all the weighted Delaunay triangulations, each τ can be represented by a rectangle $W(\tau) \times [\sigma(\tau), +\infty)$, where $W(\tau)$ is the weight interval fo τ. That is, if the weight of p is set to a value within $W(\tau)$, τ may appear and the quality of simplices incident to p is at best $\sigma(\tau)$. We take the union of the rectangles corresponding to all C_p-simplices. This produces a skyline and the highest point in the skyline yields the best weight to be assigned to p. As in sliver exudation, there is an apparent

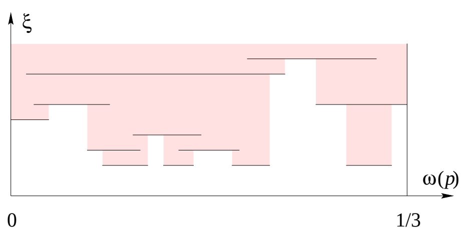

Figure 1: The skyline (borrowed from [8]).

difficulty of readmitting some j-dimensional C_p-simplex τ eliminated earlier when we process a different sample point q later. Note that $\sigma(\tau)$ is a symmetric measure, regardless whether we view τ as a C_p- or C_q-simplex. The theoretical result guarantees that there is a σ such that the weight we assign to q makes $\sigma(\tau) \geq \sigma$. Thus this readmitted C_p-simplex τ cannot be worse than what the theoretical result guarantees for p. In all there is no harm done.

If ε is sufficiently small, at the end of the pumping step, there are no j-slivers, $j \leq k+1$, in the cocones, and the last step proceeds without problem.

Running Time. Let n be the number of sample points. The time is dominated by the construction of the Voronoi diagram, which has worst-case running time $O(n^{\lceil d/2 \rceil})$.[2] The following sections show that under the sampling condition, the number of C_p-simplices incident to a sample p in all weighted Delaunay triangulations is $O(1)$, and so the running time is dominated by the first step. In the last section, we discuss how to improve the running time.

Correctness. We claim that the algorithm actually outputs $\mathrm{Del}_{\mathcal{M}}\hat{S}$ and that this is homeomorphic to the original manifold $\mathcal{M}$. The first part is verified in section 9. It depends on the success of the sliver removal step to eliminate cocone j-slivers for $j \leq k+1$, which is verified in section 8. More precisely, it is shown that (i) for σ sufficiently small, there is a weight assignment to the sample points so that no cocone j-simplex, $j \leq k+1$, is a sliver, and that (ii) for ε sufficiently small, any cocone $(k+1)$-simplex must be a sliver.

For the homeomorphism, we show that each restricted Voronoi cell is a *topological ball* [12]. The argument proceeds in three steps: (i) The normal spaces of points close on $\mathcal{M}$ are close (section 10); (ii) for any j-simplex τ, $j \leq k$, if its circumradius is $O(\varepsilon f(p))$ and neither τ nor its boundary simplices is a sliver, the normal space of τ is close to the normal space of $\mathcal{M}$ at any vertex of τ; (iii) each cell in $\mathrm{Vor}_{\mathcal{M}}\ \hat{S}$ is a topological ball (section 12). Result (ii) also shows that the output approximates $\mathcal{M}$ well in normal. With small extra effort, the distance approximation also follows.

[2] Actually, for an (ε,δ)-sampling, the worst-case size of its Delaunay triangulation may not be as large as $\Omega(n^{\lceil d/2 \rceil})$ as in the general case. This is known for dimension 3 [5].

4 Voronoi Cell Width

In this section, for each sample $p \in S$, we bound the width of $C_p \cap V_p$. We first need a technical result.

LEMMA 2. *Let y be a point inside C_p such that $\|p-y\| = c\varepsilon f(p)$ for some constant c with $c\varepsilon \leq 1/5$. Let q be the orthogonal projection of y onto $\mathcal{T}_p$. Let q' be the point on $\mathcal{M}$ closest to q. Then*

(i) $\|q-q'\| \leq c^2\varepsilon^2 f(p)/2$.

(ii) $\|q-y\| \leq c\varepsilon f(p)/10$ *and* $\|q'-y\| \leq c\varepsilon f(p)/4$.

(iii) $\|p-q'\| \geq c\varepsilon f(p)/4$.

LEMMA 3. *Assume that $\widehat{S}$ has weight property $[\omega]$. For each point $x \in C_p \cap V_p$, $\|p-x\| \leq c_1 \varepsilon f(p)$ for $c_1 = 160$ and $\varepsilon \leq 1/5c_1$.*

5 Cocone Neighbors

There are many possible weight assignments to S that has the weight property $[\omega]$. Each such weight assignment produces a possibly different set of cocone neighbors for p. We define G_p to be the set of C_p-edges that arise in *all* weight assignment with weight property $[\omega]$. In this section, we study the lengths of edges in G_p, the cardinality of G_p, and the angle between any edge in G_p and $\mathcal{T}_p$.

LEMMA 4. *Assume that $\widehat{S}$ has weight property $[\omega]$. For any edge $pq \in G_p$, $\|p-q\| \leq c_2\varepsilon f(p)$ where $c_2 = c_1(1+1/\sqrt{1-4\omega^2})$ is a constant.*

LEMMA 5. *Assume that $\widehat{S}$ has weight property $[\omega]$. Let pq and pr be edges in G_p. Then*

(i) $\|p-q\| \leq \nu \cdot \|p-r\|$ *where $\nu = c_2(\varepsilon/\delta)$ is a constant.*

(ii) $\angle(pq, \mathcal{T}_p) \leq \arcsin(c_2\varepsilon/2) \leq c_2\varepsilon$.

LEMMA 6. *For $\varepsilon > 0$ sufficiently small, there are at most λ edges in G_p, where $\lambda = (2\nu^2+1)^d$ is a constant.*

Proof. Let pr be any edge in G_p. Let u be the nearest neighbor of r. We first show that ru is an edge in G_r. It suffices to show that ru is a C_r-edge with respect to the unweighted Delaunay triangulation of S. Observe that r and u define a non-empty Voronoi facet V_{ru}. Observe that the edge ru stabs V_{ru} (let x be the midpoint of ru and suppose that q is closer to x than r and u, then $\|r-q\| \leq \|r-x\| + \|x-q\| < \|r-x\| + \|r-x\| < \|r-u\|$, which is a contradiction). By Lemma 4 and Lemma 1(ii), $\angle(ru, \mathcal{T}_r) \leq \arcsin(c_2\varepsilon/2)$. Thus, when ε is sufficiently small, ru lies inside C_r. It follows that V_{ru} intersects C_r and so ru is a C_r-edge.

We are ready to bound $|G_p|$. Without loss of generality, assume that the shortest edge in G_p has unit length. Since $\|p-r\| \geq 1$, Lemma 5(i) implies that $\|r-u\| \geq 1/\nu$. Thus if we place a ball B_r with center r and radius $1/(2\nu)$ for every edge pr in G_p, these balls are disjoint. Note that $\mathrm{vol}(B_r) = K_d(\frac{1}{2\nu})^d$, where K_d is the volume of the unit d-ball. All such B_r's lie inside a bigger ball B with center p and radius $L + 1/(2\nu)$, where L is the length of the longest edge in G_p. By Lemma 5(i), $L \leq \nu$. Therefore, $\mathrm{vol}(B) \leq K_d(\frac{2\nu^2+1}{2\nu})^d = (2\nu^2+1)^d \mathrm{vol}(B_r)$. Hence there are at most $(2\nu^2+1)^d$ edges in G_p. ⊔⊓

6 Orthoradius-Edge Ratio

In this section, we bound the orthoradius-edge ratio of C_p-simplices. The orthoradius-edge ratio of a simplex τ is R'_τ/L_τ, where R'_τ is the radius of the smallest orthosphere of τ (its center lies in $\mathrm{aff}(\tau)$), and L_τ is the shortest edge length of τ. Note that the smallest orthosphere of τ is not necessarily further than orthogonal from other weighted vertices in $\widehat{S}$. Also note that the circumradius-edge and orthoradius-edge ratios may be quite different.

LEMMA 7. *Assume that $\widehat{S}$ has weight property $[\omega]$. For any C_p-simplex τ, $R'_\tau \leq \rho L_\tau$, where $\rho = 5\nu(\varepsilon/\delta)$ is a constant.*

7 Weight Interval

When we pump a vertex p of a sliver τ, τ may remain a C_p-simplex for a while. In this section, we bound the length of the weight (radius) interval for p such that τ remains a C_p-simplex. We first prove two technical results.

LEMMA 8. *Assume that $\widehat{S}$ has weight property $[\omega]$. Let $\tau = p * \tau_p$ be a C_p-simplex. The distance between the orthocenter of τ and $\mathrm{aff}(\tau_p)$ is at most $c_3\varepsilon f(p)$ for the constant $c_3 = c_1 + c_2(1+\omega+\rho\nu)$.*

For a simplex τ and a vertex q of τ, let D_q be the distance from q to $\mathrm{aff}(\tau_q)$ (recall $\tau = q * \tau_q$).

LEMMA 9. *Assume that $\widehat{S}$ has weight property $[\omega]$. Let τ be a j-dimensional C_p-simplex. If τ is a sliver, then $D_q < j\sigma L_\tau$, for some vertex q of τ. If neither τ nor its boundary simplices are slivers, then $D_q \geq j\sigma L_\tau$ for each vertex q.*

We are ready to bound the weight interval.

LEMMA 10. *Let τ be a j-dimensional C_p-simplex. If τ is a sliver, τ remains a C_p-simplex in an interval of squared weight that has length at most $4c_2c_3j\sigma\varepsilon^2 f(p)^2$.*

Proof. Let p be a vertex of τ such that $D_p < j\sigma L_\tau$. Let $\xi(P^2)$ be the signed distance of the orthocenter z of τ from $\mathrm{aff}(\tau_p)$ when p has squared weight P^2 (see figure). $\xi(P^2)$ is positive if the orthocenter of τ and p lie on the same side of $\mathrm{aff}(\tau_p)$; otherwise $\xi(P^2)$ is negative. It has been proved

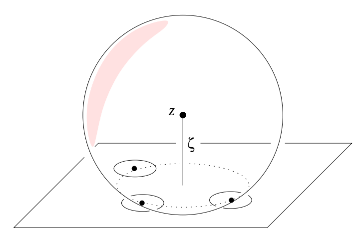

Figure 2: The orthosphere of three weighted points on the plane shown and a fourth one off but close to the plane. The distance of the orthocenter z to the plane is very sensitive to change in the weight of the fourth point (borrowed from [8]).

in [8] that

$$\xi(P^2) = \xi(0) - \frac{P^2}{2D_p}.$$

By Lemma 8, for τ to be a C_p-simplex, $|\xi(P^2)| \leq c_3\varepsilon f(p)$. Substituting into the above, we get

$$2D_p(\xi(0) - c_3\varepsilon f(p)) \leq P^2 \leq 2D_p(\xi(0) + c_3\varepsilon f(p)).$$

This implies that the length of the squared weight interval is at most $4c_3\varepsilon D_p f(p)$. Using Lemma 9 and Lemma 4, we get $4c_3\varepsilon D_p f(p) \leq 4c_3 j\sigma\varepsilon L_\tau f(p) \leq 4c_2c_3j\sigma\varepsilon^2 f(p)^2$. ⊔⊓

8 Pumping Slivers

Our strategy is to pump slivers in increasing order of their dimensions. That is, we first pump all 3-slivers, then 4-slivers, and so on. Therefore, when we consider a j-simplex τ, all the boundary simplices of τ are guaranteed not to be slivers. This will allow us to invoke Lemma 10.

Consider a sample p. By Lemma 6 and Lemma 10, $j\lambda^j \cdot 4c_2c_3\sigma\varepsilon^2 f(p)^2$ is an upper bound on the total length of the squared weight intervals for the j-slivers to remain C_p-simplices. Summing over the range $[3, k+1]$ for j, we obtain an upper bound of

$$(k+1)\lambda^{k+2} \cdot 4c_2c_3\sigma\varepsilon^2 f(p)^2.$$

By the (ε, δ)-sampling, the maximum weight of p is at least $\omega^2\delta^2 f(p)^2$. Therefore, if we choose σ such that

$$\sigma < \frac{1}{(k+1)\lambda^{k+2}} \cdot \frac{\omega^2}{4c_2c_3} \cdot \left(\frac{\delta}{\varepsilon}\right)^2,$$

then we can assign a weight to p such that no C_p-simplex is a sliver. (Recall that δ/ε is a constant.) If k is not known, we can replace k by d and the resulting choice of σ is guaranteed to work, albeit even more pessimistic.

The following lemmas show that, if ε is sufficiently small, provided its boundary simplices are not slivers, a $(k+1)$-dimensional C_p-simplex is a sliver. Hence after pumping, no C_p-simplex has dimension $k+1$ or higher. We first bound the circumradius-edge ratio of a non-sliver.

LEMMA 11. *Let τ be a j-dimensional C_p-simplex. If τ and its boundary simplices are not slivers, its circumradius-edge ratio is bounded by the constant $\gamma_j = (\nu^2/\sigma)^{2^j-1}$.*

Proof. We prove by induction on j. For $j = 1$, τ is an edge. We define the circumradius of an edge to be half of its length. Then the circumradius-edge ratio is 1/2 which is less than $\gamma_1 = \nu^2/\sigma$. Assume that $j > 1$. Let z be the circumcenter and p a vertex of τ. Recall that $\tau = p * \tau_p$. There are two cases to consider. Let R be the circumradius of τ.

Case 1: $z \in \mathrm{int}(\tau)$. Let H be a $(j-1)$-flat in $\mathrm{aff}(\tau)$ that passes through z and is orthogonal to pz. Since $z \in \mathrm{int}(\tau)$, H separates a vertex q of τ from p. It follows that $\angle pzq > \pi/2$ and $\|p-q\| \geq R$. Thus $R/L_\tau \leq \|p-q\|/L_\tau \leq \nu$, by Lemma 5(i), which is less than γ_j.

Case 2: $z \notin \mathrm{int}(\tau)$. Let az be the radius of the circumsphere of τ such that az is orthogonal to $\mathrm{aff}(\tau_p)$. Let x be the point $az \cap \mathrm{aff}(\tau_p)$. Let q be a vertex of τ_p. Let R' be the circumradius of τ_p. Since $z \notin \mathrm{int}(\tau)$, we have

$$\|a-q\| \leq \sqrt{2}\,R' \leq \sqrt{2}\gamma_{j-1}L_{\tau_p}$$

by induction assumption. Observe that $\sin\angle aqx = \|a-q\|/(2R)$. Using the previous inequality and Lemma 5(i), we get

$$\sin\angle aqx \leq \frac{\sqrt{2}\gamma_{j-1}L_{\tau_p}}{2R} \leq \frac{\gamma_{j-1}\nu L_\tau}{\sqrt{2}R}.$$

The volume of τ is at most $\frac{1}{j} \cdot \mathrm{vol}(\tau_p) \cdot \|a-q\| \cdot \sin\angle aqx$. Let L be the maximum edge length of τ. Using the previous inequalities for $\|a-q\|$ and $\sin\angle aqx$, we obtain

$$\mathrm{vol}(\tau) \leq \frac{\mathrm{vol}(\tau_p)}{j} \cdot \sqrt{2}\gamma_{j-1}L \cdot \frac{\gamma_{j-1}\nu L_\tau}{\sqrt{2}R} \leq \mathrm{vol}(\tau_p)\frac{\gamma_{j-1}^2\nu^2 L_\tau^2}{R}.$$

Since τ is not a sliver, $\mathrm{vol}(\tau) \geq \sigma L_\tau \mathrm{vol}(\tau_p)$. Substituting above, we get $R/L_\tau \leq \gamma_{j-1}^2(\nu^2/\sigma)$. Thus $R/L_\tau \leq \gamma_j$ where $\gamma_j = \gamma_{j-1}^2(\nu^2/\sigma)$. $\square$

LEMMA 12. *Let τ be a j-dimensional C_p-simplex. Let x be a point in τ. If τ and its boundary simplices are not slivers, then $\angle(px, \mathcal{T}_p) \leq 4c_2\gamma_j\varepsilon$.*

Proof. Take a d-dimensional medial sphere that touch $\mathcal{M}$ at p. Shrink this sphere towards p until its radius becomes $f(p)$. Denote the resulting sphere by M_1. Let M_2 be another sphere with radius $f(p)$ such that M_2 touches p at $\mathcal{M}$ and p is the midpoint of the centers of M_1 and M_2. The smallest circumsphere of τ intersects M_1 and M_2 at two hyperspheres C_1 and C_2, respectively. By Lemma 4 and Lemma 11, the circumradius of τ is at most $c_2\gamma_j\varepsilon f(p)$. Thus the angle between the normal of the support hyperplanes of C_1 and the vector from p to the center of M_1 is at most $\arcsin(c_2\gamma_j\varepsilon)$. The same holds for the normal of the support hyperplane of C_2. It follows that the support hyperplanes of C_1 and C_2 make a wedge of angle at most $2\arcsin(c_2\gamma_j\varepsilon)$. Since the vertices of τ lie outside M_1 and M_2, they lie within this wedge. This implies that px lie within the wedge too. Since $\mathcal{T}_p$ cuts through the wedge, $\angle(px, \mathcal{T}_p) \leq 2\arcsin(c_2\gamma_j\varepsilon)$, which is at most $4c_2\gamma_j\varepsilon$ for sufficiently small ε. $\square$

LEMMA 13. *Let $k = \dim(\mathcal{M})$. Assume that $k \geq 2$, $\varepsilon < (k+1)\sigma/(1+4\gamma_k)\nu$, and $\widehat{S}$ has weight property $[\omega]$. Let τ be a $(k+1)$-dimensional C_p-simplex. If the boundary simplices of τ are not slivers, τ is a sliver.*

Proof. Let τ be a $(k+1)$-dimensional C_p-simplex. Recall that $k + \dim(\mathcal{N}_p)$ is equal to the dimension of the underlying space. Thus, there is some unit normal $\vec{n} \in \mathcal{N}_p$ such that $p + \vec{n} \in \mathrm{aff}(\tau)$. Without loss of generality, we treat $\vec{n}$ as the vertical axis of $\mathrm{aff}(\tau)$.

For each vertex r of τ other than p, let $\tau = r * \tau_r$, as usual, and let r' be the projection along $\vec{n}$ of r onto $\mathrm{aff}(\tau_r)$. We claim that there is a vertex $q \neq p$ of τ such that the support line of pq' intersects τ_q at a point other than p. There are two cases.

Case 1: there is a vertical k-flat H in $\mathrm{aff}(\tau)$ through p and containing $\vec{n}$ such that at least three other vertices of τ lie on one side H^+ of H. Rotating H around $\vec{n}$ brings it into contact with two vertices a and b of τ in H^+. Let q be any vertex of τ in H^+ other than a and b. The orthogonal projection of pq onto the plane of abp intersects abp at a point other than p. It follows that pq' intersects τ_q at a point other than p.

Case 2: the k-flat in case 1 does not exist. Let H be any k-flat in $\mathrm{aff}(\tau)$ through p and containing $\vec{n}$. Since $k \geq 2$, there must be exactly two vertices of τ on one side H^+ of H. Let these two vertices be denoted by a and b. Let H^- denote the side opposite to H^+. If we extend ap and bp into H^-, we obtain a 2-d cone C in the plane of abp in H^-. For any vertex q of τ in H^-, the projection of pq onto the plane of abp must lie inside the 2-d cone C; otherwise, there would be a k-flat that have a, b, and q on the same side, contradicting the assumption that case 1 does not apply. Thus, the support line of the projection of pq onto the plane of abp intersects abp at a point other than p. It follows that the support line of pq' intersects τ_q at a point other than p.

This completes the proof of our claim. Now, let $x \neq p$ be a point in the intersection of the support line of pq' and τ_q. By Lemma 12, applied to τ_q, $\angle(pq', \mathcal{T}_p) \leq 4c_2\gamma_k\varepsilon$. By Lemma 4 and Lemma 1(ii), $\angle(pq, \mathcal{T}_p) \leq \arcsin(c_2\varepsilon/2)$, which is at most $c_2\varepsilon$ for sufficiently small ε. As qq' is parallel to $\vec{n}$, we conclude that $\angle qpq' \leq \angle(pq, \mathcal{T}_p) + \angle(pq', \mathcal{T}_p) \leq c_2(1 + 4\gamma_k)\varepsilon$.

The height of q from $\mathrm{aff}(\tau_q)$ is at most $\|p-q\| \cdot \sin\angle qpq' \leq \|p - q\| \cdot \sin(c_2(1 + 4\gamma_k)\varepsilon) \leq c_2(1 + 4\gamma_k)\varepsilon L$, where L is the maximum edge length of τ. Thus

$$\mathrm{vol}(\tau) \leq \frac{c_2(1 + 4\gamma_k)\varepsilon L}{k + 1}\mathrm{vol}(\tau_q).$$

By Lemma 5(i), $L \leq \nu L_\tau$. Thus, if

$$\varepsilon < \frac{(k + 1)\sigma}{(1 + 4\gamma_k)\nu},$$

then τ is a sliver. ◻

9 Algorithm Output

We show that our algorithm actually outputs $\mathrm{Del}_{\mathcal{M}}(\hat{S})$. Let X be the set of all simplices output by our reconstruction algorithm. Recall that X is the set of cocone simplices.

LEMMA 14. *X is* $\mathrm{Del}_{\mathcal{M}}(\hat{S})$.

Proof. The assignment of weights of our algorithm ensures that no Voronoi cell of dimension less than $d - k$ intersects the cocone of a point p in S. So certainly, no Delaunay simplex of dimension larger than k is in X. Also certainly, any simplex τ in $\mathrm{Del}_{\mathcal{M}}(\hat{S})$ is in X because by definition its dual Voronoi cell V_τ intersects $\mathcal{M}$ and hence V_τ intersects the cocone of the vertices of τ. It remains to see that there is no Voronoi cell V_τ that intersects the cocone C_p of a vertex p of its dual simplex τ, but it does not intersect $\mathcal{M}$. For the sake of contradiction, let V_τ be such a Voronoi cell of smallest dimensionality and let x be a point of V_τ inside C_p. Also, let T be the intersection of $\mathcal{T}_p$ with $\mathrm{aff}(V_\tau)$. Inside $\mathrm{aff}(V_\tau)$, let N be the orthogonal complement of T through x. N must intersect $\mathcal{M}$ inside C_p. Since V_τ does not intersect $\mathcal{M}$ then N should intersect inside C_p a smaller dimensional Voronoi cell that bounds V_τ and which also does not intersect $\mathcal{M}$. This is a contradiction. ◻

Section 11 shows that $\mathrm{Del}_{\mathcal{M}}(\hat{S})$ approximates $\mathcal{M}$ well in normal, and section 12 shows that $\mathrm{Del}_{\mathcal{M}}(\hat{S})$ is homeomorphic to $\mathcal{M}$.

10 Normal Variation

The proof of the following lemma extends that of a 3-d version that appears in [1].

LEMMA 15. *Let $p, q \in \mathcal{M}$ such that $\|p-q\| \leq c\varepsilon f(p)$, then $\angle\mathcal{N}_p\mathcal{N}_q \leq c_4 c\varepsilon$ for some constant c_4.*

Proof. Consider the line segment pq joining p and q and let $p(t)$ be a linear parametrization of pq in the interval $[0, 1]$. For $t \in [0, 1]$, let $g(t)$ be the closest point to $p(t)$ in $\mathcal{M}$. Since pq is away from the medial axis, $g(t)$ is well-defined (there is a unique closest point) and also one-to-one (if x is closest for p' and p'' in pq then both $p'x$ and $p''x$ are normal to $\mathcal{M}$ at x and hence pq is in the normal space of $\mathcal{M}$ at x; therefore the diametral sphere of px is tangent to $\mathcal{M}$ at x, and so $\|p - x\| \geq 2f(x)$, which is in contradiction with $\|p - q\| \leq c\varepsilon f(p)$ for c and ε sufficiently small). The function $g(t)$ is indeed smooth. Let γ be the curve in $\mathcal{M}$ described by $g(t)$, and let dt and ds be the lengths of corresponding infinitesimal segments on pq and γ, that is, $ds = \|g(t + dt) - g(t)\|$. We claim that $ds \leq 4dt$. To see this, first consider the medial ball B tangent to $\mathcal{M}$ at $g(t)$ and with center on the ray from $g(t)$ towards $p(t)$. The radius of B is greater than $f(g(t))$ and so greater than $c' f(p)$ for some constant c' (by Lipschitz property of f), and also the ball B' centered at $p(t+dt)$ and passing through $g(t)$. Note that $g(t+dt)$ must lie in the portion of B' outside B (since B is a medial ball and hence its interior is disjoint from $\mathcal{M}$, and since $g(t+dt)$ cannot be further from $p(t+dt)$ than $g(t)$). This portion lies within distance $4dt\sin\theta$ from $g(t)$ where θ is the angle between pq and $p(t)g(t)$: Consider the figure, which shows the 2-flat spanned by pq and $p(t)g(t)$, where $p' = p(t)$, $p'' = p(t+dt)$, $dt = \|p' - p''\|$, $q' = g(t)$, z is the center and R the radius of B, q'' is the projection of q' on the line that contains zp'', $d = \|p' - q'\|$ and $h = \|q' - q''\|$. An elementary calculation shows that

$$h \leq \frac{R}{R-d} \cdot dt \sin\theta$$

which is smaller than $2dt\sin\theta$ for ε sufficiently small so that $d \leq R/2$. Finally, $g(t+dt)$ lies within distance $2h$ from $q = g(t)$, that is, within distance $4dt\sin\theta$. This is at most $4dt$.

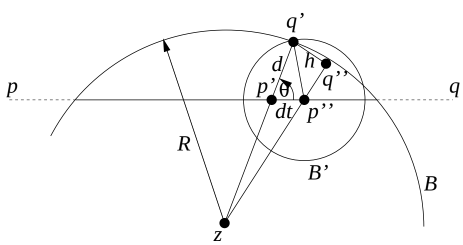

Figure 3: The point on $\mathcal{M}$ closest to $p'' = p(t+dt)$ must be closer than the closest point $q' = g(t)$ to $p' = p(t)$ and outside of the medial ball B of $\mathcal{M}$ at q' with center on the ray from q' towards p'.

Now, for a unit normal $\vec{n} \in \mathcal{N}_p$, let $\vec{n}(t)$ be the unit normal in $\mathcal{N}_{g(t)}$ that forms a smallest angle with $\vec{n}$. Thus, $\vec{n}(t)$ is the normalized projection of $\vec{n}$ on $\mathcal{N}_{g(t)}$ and $\vec{n}(0) = \vec{n}$. We claim that $\angle\vec{n}(t)\vec{n}(t+dt)$ is bounded by $(4/c'f(p))ds$. To verify this, consider the set of balls $B_{\vec{n}'}$, with radius $R = f(g(t))$ and tangent to $\mathcal{M}$ at $g(t)$ in the direction $\vec{n}'$, where $\vec{n}'$ is a normal direction at $g(t)$. Because of the balls $B_{\vec{n}'}$, the rate of change of the normal to $\mathcal{M}$ in any direction with respect to $ds = \|g(t+dt) - g(t)\|$ is bounded by $1/R$. So $\angle\vec{n}(t)\vec{n}(t+dt) \leq ds/R$. Since $R = f(g(t))$ is at least $c'f(p)$, then $\angle\vec{n}(t)\vec{n}(t+dt)$ is upper bounded by $ds/c'f(p)$, which is at most $4dt/c'f(p)$ by the argument above. Adding this bound over $[0,1]$, we obtain that $\angle\vec{n}(0)\vec{n}(1)$ is at most $4\|\|p-q\|/c'f(p)$, which is at most $c_4 c\varepsilon$ for some constant c_4. ◻

11 Normal Approximation

The conditions on the simplex τ in the following lemma hold for the simplices computed by our algorithm. The lemma implies that the reconstruction produced approximates $\mathcal{M}$ well in the sense of normal approximation. This result is also useful in proving that the restricted Voronoi cells are topological balls. In the following proof, the term cocone refers to the complement of a (usual) double cone around an specified direction; its aperture is $\pi/2$ minus the aperture of the cone. For a simplex τ, a vector $\vec{n}_\tau$ is normal to τ if it is orthogonal to $\mathrm{aff}(\tau)$.

LEMMA 16. *Suppose τ is a j-simplex for $j \leq k$, with vertices on $\mathcal{M}$, circumradius $O(\varepsilon f(p))$, where p is one of its vertices, and such that neither τ nor its boundary simplices is a sliver. Then for any normal $\vec{n}_p$ of $\mathcal{M}$ at p, τ has a normal $\vec{n}_\tau$ such that $\angle\vec{n}_p\vec{n}_\tau$ is at most $k_j\varepsilon$, for some constant k_j.*

Proof. The proof is by induction on j. For $j = 0$, the claim is trivial. For $j \geq 1$, let $\tau = q * \tau_q$ as usual and let D_q be the distance from q to $\mathrm{aff}(\tau_q)$. Because τ is not a sliver, $D_q \geq j\sigma L_\tau \geq d_j\varepsilon f(p)$ with d_j a constant that depends on j and the dimension d. By induction, τ_q has a normal $\vec{n}_{\tau_q}$ such that $\angle\vec{n}_p\vec{n}_{\tau'}$ is at most $k_{j-1}\varepsilon$. Let q' be the point in $\mathrm{aff}(\tau_q)$ closest to q, let h be the $(d-1)$-flat (hyperplane) containing τ_q and normal to $\vec{n}_{\tau_q}$, and let γ be the $(d-2)$-flat in h orthogonal to qq'. Consider now rotating h around γ until a hyperplane h' that contains q is obtained; note that its normal is also a normal of τ, and so we denote it with $\vec{n}_\tau$. We claim that $\angle\vec{n}_{\tau_q}\vec{n}_\tau$ is at most $(c'/d_j)k_{j-1}\varepsilon$ for some constant c'. This will imply that $\angle\vec{n}\vec{n}_\tau$ is at most $k_{j-1}\varepsilon + (c'/d_j)k_{j-1}\varepsilon$ and so at most $k_j\varepsilon$ where k_j is a solution to the recurrence $k_j = k_{j-1}(1 + c'/d_j)$. To complete the proof, we verify the claim as follows. q is in a cocone C_q around $\vec{n}_{\tau_q}$ of aperture $2k_{j-1}\varepsilon$: it is in a cocone around $\vec{n}_p$ of aperture $c\varepsilon$ by the dense sampling, and since we are changing the reference to $\vec{n}_{\tau_q}$, then we need to increase the aperture by $k_{j-1}\varepsilon$ (we set $c_0 = c$ so that $c \leq k_{j-1}$). We want to see that q is also in a *wedge* ω around γ, obtained by pivoting h around γ, with an angle α that is at most $(c/d_j)k_{j-1}\varepsilon$. An elementary geometric computation shows that if q is in the cocone C_p with aperture φ and at distance R from p, and D_q from γ, then it is inside the wedge ω of angle α, if $\sin\alpha \leq (R/D_q)\sin\varphi$ (see figure). Since $R = \|p-q\| \leq c_2\varepsilon f(p)$, $D_q \geq d_j\varepsilon f(p)$, and $\varphi \leq 2k_{j-1}\varepsilon$, then for $\sin\alpha \leq (c_2/d_j)\sin(2k_{j-1}\varepsilon)$, q is in the wedge ω around γ. This means that $\angle\vec{n}_\tau\vec{n}_{\tau_q} \leq (c'/d_j)k_{j-1}\varepsilon$, for some constant c', as we had claimed. ◻

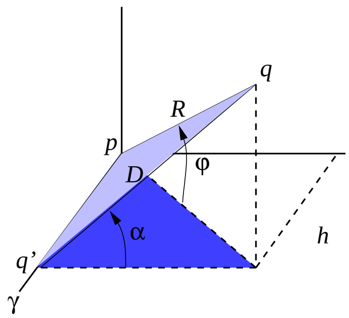

Figure 4: q is exactly on the boundary of a cocone of p around $\vec{n}_{\tau_q}$ (the vertical) with aperture φ and a wedge around γ of aperture α.

Using this lemma, we choose ε sufficiently small so that each simplex in $\mathrm{Del}_{\mathcal{M}}\hat{S}$ has a normal within say $\pi/32$ of a normal of either of its vertices. We call this *good normal approximation* for the simplices.

Remark. Removing slivers is an essential part of our reconstruction algorithm and its proof of correctness. In a way, this is actually needed if we want to guarantee good normal approximation for the simplices. The normal of a sliver can be arbitrarily wrong even if its circumradius is $O(\varepsilon f(p))$ and the circumradius-edge ratio is bounded. Consider a cube in $\mathbb{R}^4$ with side length d and smooth out its ridges and corners, to get a smooth 3-manifold $\mathcal{M}$. Close to the center of the facets, the manifold is flat and the local feature size is $\Theta(d)$. Consider the Delaunay triangulation of a dense sampling on $\mathcal{M}$ such that the circumradius-edge radius for every tetrahedron is larger than a constant. In the central portion of a facet, locally, this is a 3-d triangulation (since the manifold is flat there), and the normal of all simplices are correct (point to the 4th dimension). Suppose there is a sliver there, made up of a triangle qrs and an extra vertex p, so that p is at a distance Δ from $\mathrm{aff}(qrs)$ with Δ arbitrarily small (this is possible under the condition of a dense sampling and a triangulation with bounded circumradius-edge ratio). To be precise, let $p = (0, 0, \Delta, 0)$, $q = (1, 0, 0, 0)$, $r = (1, 1, 0, 0)$ and $s = (0, 1, 0, 0)$. Then the normal is in the direction $(0, 0, 0, 1)$. Now, we deform very slightly $\mathcal{M}$ near $((0, 0, 0, 0)$ into the 4th dimension –creating a very small *bump*– (recall $\mathcal{M}$ was flat there) to obtain a manifold $\mathcal{M}'$, and move p into the 4th dimension into the bump, also a distance Δ away from $\mathrm{aff}(qrs)$. More precisely $p' = (0, 0, 0, \Delta)$. Since Δ is very small, this can be done without changing significantly the local feature size of $\mathcal{M}$ and, in particular, even near p', the local feature size remains essentially the same (say the bump has curvature radius $\Theta(d)$). However, the normal is in the direction $(0, 0, 1, 0)$. Thus, after moving p, we still have an ε-sampling for the new manifold $\mathcal{M}'$, and also the restricted Delaunay triangualtion remains the same except for the slightly changed sliver $pqrs$ and its neighbors. We actually want the new $pqrs$ to be in the restricted Delaunay triangulation; to achieve this, we actually need to be more careful in moving p: move p so that there is still a restricted Voronoi vertex corresponding to the sliver on the manifold; this can be achieved by moving p on the circumsphere of q, r, s.

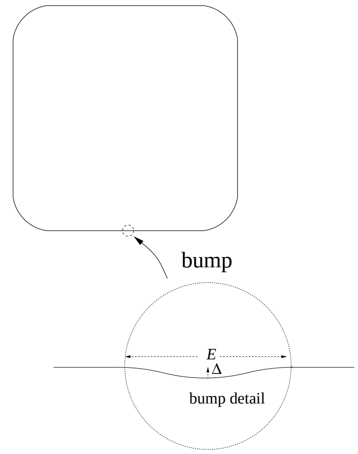

Figure 5: A small bump is introduced in the cube manifold, detailed on the right. The effect of the bump on the local feature size at any point is negligible (even at points on the bump). For any $\varepsilon > 0$, E can be chosen sufficiently small and Δ even smaller, so that for the corresponding manifold with bump, there is an ε-sampling in which p is the only sample point in a large neighborhood of the bump.

12 Ball Property for Cells in $\mathrm{Vor}_{\mathcal{M}}\ \hat{S}$

To verify that $\mathrm{Del}_{\mathcal{M}}(\hat{S})$ is homeomorphic to $\mathcal{M}$, it suffices to show that each of the cells in $\mathrm{Vor}_{\mathcal{M}}(\hat{S})$ is homeomorphic to a ball [12] (the result there is proved for unweighted points but carries over to the weighted case given the weight property).

We assume in this section that ε is sufficiently small and that no simplex dual to a restricted Voronoi cell is a sliver, so that good normal approximation for these simplices hold, say with angle $\pi/32$.

We need the following lemma, which is a minor modification of lemma 3.

LEMMA 17. *If $x, y \in \mathcal{M}$ belong to a common cell of* $\mathrm{Vor}\ \hat{S}$, *then* $\|x - y\| \leq c_5 \varepsilon f(x)$.

LEMMA 18. *For (ε, δ)-sampling with ε sufficiently small, and assuming that no restricted Voronoi cell is dual to a sliver, then each j-cell of* $\mathrm{Vor}_{\mathcal{M}}\ \hat{S}$ *is a topological ball.*

13 Improvements

Improved Running Time. As described in Section 3, our algorithm requires the computation of the complete (weigthed) Delaunay/Voronoi complex. Though the complete complex is important for the determination of the poles and the cocone, in the end, for the output complex only the simplices that connect each sample point with other sample points in a small neighborhood are needed, and they do not depend on other distant samples.

Giesen and Wagner [15] have shown that, under (ε, δ)-sampling, the manifold dimension can be estimated from an appropriate neighborhood of each point. For a sample $p \in S$, its α-*neighborhood* is

$$N_\alpha(p) = \{q \in S - \{p\} : \|p - q\| \leq \alpha \min_{q' \in S - \{p\}} \|p - q'\|\}.$$

For $\alpha \approx 2\varepsilon/(1-\varepsilon)\delta$, $N_\alpha(p)$ has size $O(1)$ and captures the shape locally: they fit an l-dimensional flat to the set $N_\alpha(p)$ with $l = 1, \ldots, d$; the fitting error is larger than a threshold if $l < k$, and smaller if $l \geq k$, and so k can be determined. The fitting k-flat is then a good approximation for $\mathcal{T}_p$. These neighborhoods can also be used as the basis for steps 3 and 4 of the algorithm in Section 3: the (weighted) Delaunay simplices incident to a sample p that are relevant for our reconstruction algorithm can be determined from $N_\alpha(p)$ (for α slightly larger). Using approximate nearest neighbor reporting [4], $N_\alpha(p)$ can be computed for all $p \in S$ in time $O(n \log n)$. Then dimension detection [15] and sliver exudation take time $O(n)$. We are hiding large constants depending on the dimension in all cases. The overall running time is $O(n \log n)$.

Weaker Sampling. (ε, δ)-sampling is somewhat restrictive. Our approach applies to a more relaxed variant: The parameter ε does not need to be the same throughout the manifold; it suffices that there is an ε locally that changes slowly over the manifold –this is called *locally uniform* sampling in [13]–, while the ratio ε/δ remains constant. This includes, for example, a globally uniform sample, case not included in (ε, δ)-sampling. The algorithm remains the same; the proof of correctness extends without problem.

The result can potentially be extended to (ϵ, ℓ)-sampling, at the cost of even denser sampling. In this case, we would need a preprocessing step that enforces locally uniform sampling, similar to that in [13]. Briefly, it consists in *decimating* the sample set to enforce the local uniformity. For this a local decimation radius need to be estimated. Using approximation techniques, all this can be done in time $O(n \log n)$. If the manifold dimension k is known, then the extension could be even to ε-sampling (no lower bound on distances between samples). Details for these last extensions still need to be worked out.

References

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

[2] N. Amenta, M. Bern and D. Eppstein. The crust and the β-skeleton: combinatorial curve reconstruction. *Graphical Models and Image Processing*, **60** (1998), 125-135.

[3] N. Amenta, S. Choi, T. K. Dey and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. *Internat. J. Comput. Geom. Applications*, **12** (2002), 125–141.

[4] S. Arya, D. M. Mount, N. S. Netanyahu and R. Silverman. An optimal algorithm for approximate nearest neighbor searching in fixed dimension. *J. ACM* 45(6):891–923, (1998).

[5] D. Attali, J.-D. Boissonnat and A. Lieutier. Complexity of the delaunay triangulation of points on surfaces: the smooth case. In *Proc. 19th Ann. Sympos. Comput. Geom.*, (2003), 201–210.

[6] C. Bregler and S. M. Omohundro. Nonlinear manifold learning for visual speech recognition. *Proc. 5th Internat. Conf. Comput. Vision*, (1995), 494-499.

[7] J. D. Boissonnat and F. Cazals. Smooth surface reconstruction via natural neighbor interpolation of distance functions. *Proc. 16th Ann. Sympos. Comput. Geom.*, (2000), 223–232.

[8] S.-W. Cheng, T.K. Dey, H. Edelsbrunner, M.A. Facello, and S.-H. Teng. Sliver Exudation. *Journal of ACM*, 47 (2000), 883–904.

[9] T. F. Cox and M. A. A. Cox. *Multidimensional scaling*. Chapman and Hall, London, 1994.

[10] T. K. Dey and P. Kumar. A simple provable algorithm for curve reconstruction. *Proc. 12th Ann. ACM-SIAM Sympos. Discrete Algorithms*, (1999), 893–894.

[11] T. K. Dey, J. Giesen, S. Goswami and W. Zhao. Shape dimension and approximation from samples. *Discrete Comput. Geom.* **29** (2003), 419–434.

[12] H. Edelsbrunner and N. Shah. Triangulating topological spaces. *Proc. 10th ACM Sympos. Comput. Geom.*, (1994), 285–292.

[13] S. Funke and E.A. Ramos. Smoooth-surface reconnstruction in near-linear time. In *Proc. 13th ACM-SIAM Symposium on Discrete Algorithms,*, (2002).

[14] J. Giesen. Curve reconstruction, the TSP, and Menger's theorem on length. *Proc. 15th Ann. Sympos. Comput. Geom.*, (1999), 207–216.

[15] J. Giesen and U. Wagner. Shape dimension and intrinsic metric from samples of manifolds. *Proceedings of the 19th Annual ACM Symposium on Computational Geometry*, 2003, 329–337.

[16] P. Indyk. Algorithmic applications of low-distorion geometric embeddings. *42nd Annu. IEEE Sympos. Foundations Comput. Sci.*, (2001), 10–33.

[17] T. Martinetz and K. Schulten. Topology preserving networks. *Neural Networks*, **7** (1994), 507–522.

[18] J. B. Tenenbaum, V. de Silva and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science* **290** (5500), 2000.

[19] K. M.-K. Yip. *KAM : A System for Intelligently Guiding Numerical Experimentation by Computer*. The MIT Press, Cambridge, Massachusetts, 1991.

A Omitted Proofs

A.1 Proof of Lemma 2 *Proof.* Since $\|p-q\| \le \|p-y\| = c\varepsilon f(p)/2$, Lemma 1(iii) implies that $\|q-q'\| < c^2\varepsilon^2 f(p)/2$. This proves (i). Since py lies inside C_p, $\|q-y\| \le \|p-y\| \cdot \sin(\pi/16)$, where $\pi/16$ accounts for half of the angular aperture of C_p and the approximation error of $\mathcal{T}_p$ through $\mathrm{aff}(V_p^k)$. Thus,

$$\|q-y\| \le \frac{c\varepsilon}{2} f(p) \cdot \sin\frac{\pi}{16} \le \frac{c\varepsilon}{10} f(p).$$

Thus triangle inequality implies that

$$\begin{aligned} \|q'-y\| &\le \|q-y\| + \|q-q'\| \\ &< \left(\frac{c\varepsilon}{10} + \frac{c^2\varepsilon^2}{2}\right) f(p), \end{aligned}$$

which is at most $c\varepsilon f(p)/4$ for sufficiently small ε. This proves (ii). By triangle equality,

$$\begin{aligned} \|p-q'\| &\ge \|p-y\| - \|q'-y\| \\ &\ge \frac{c\varepsilon}{2} f(p) - \frac{c\varepsilon}{10} f(p) - \frac{c^2\varepsilon^2}{2} f(p) \\ &\ge \frac{c\varepsilon}{4} f(p), \end{aligned}$$

since $\varepsilon c \le 1/5$. ◻

A.2 Proof of Lemma 3 *Proof.* Let B be the ball centered at p with radius $c_1\varepsilon f(p)$. Assume that the contrary that there is a point $x \in C_p \cap V_p - B$. We use $cone(px)$ to denote the cone with axis px, apex p, and aperture $\pi/3$. Let y be the point on px such that $\|p-y\| = c_1\varepsilon f(p)/2$. Let q be the orthogonal projection of y onto $\mathcal{T}_p$. Let q' be the point on $\mathcal{M}$ closest to q. By Lemma 2(ii),

$$\sin\angle ypq = \frac{\|q-y\|}{\|p-y\|} \le \frac{1}{5} \quad \Rightarrow \quad \angle ypq \le \frac{\pi}{15}.$$

By Lemma 2(i), $\|q-q'\| \le c_1^2\varepsilon^2 f(p)/8$. By considering the triangle pqq' and using Lemma 2(i) and (iii), we have

$$\sin\angle qpq' \le \frac{\|q-q'\|}{\|p-q'\|} \le c_1\varepsilon.$$

It follows that $\angle qpq' \le 2c_1\varepsilon$ for sufficiently small ε. We have $\angle ypq' \le \angle ypq + \angle qpq' \le \pi/15 + 2c_1\varepsilon$. We conclude that pq' lies inside $cone(px)$, and the angle between pq' and the boundary of $cone(px)$ is at least $\pi/6 - \pi/15 - 2c_1\varepsilon = \pi/10 - 2c_1\varepsilon$, which is at least $\pi/20$ for sufficiently small ε.

Note that $\|p-q'\| \le \|p-y\| + \|y-q'\| \le c_1\varepsilon f(p)/2 + c_1\varepsilon f(p)/4 = (3c_1/4)\varepsilon f(p) \le f(p)$ for $\varepsilon \le 4/3c_1$, where we have used Lemma 2(ii). By Lipschitz property, $f(q') \le f(p) + \|p-q'\| \le 2f(p)$. Let B' be the ball centered at q' with radius $\|p-q'\| \cdot \sin(\pi/20)$. By Lemma 2(iii), the radius of B' is at least $(c_1\varepsilon f(p)/8) \cdot \sin(\pi/20) \ge (c_1\varepsilon f(p)/8) \cdot 0.1 \ge 2\varepsilon f(p) \ge \varepsilon f(q')$ as $c_1 \ge 160$. Note that B' does not contain p, and B' lies inside B. By the (ε,δ)-sampling, B' contains a sample r. Then

$$\begin{aligned} \pi_p(x) - \pi_r(x) &= \|p-x\|^2 - \|r-x\|^2 - P^2 + R^2 \\ &\ge \|p-x\|^2 - \|r-x\|^2 - P^2 \\ &\ge \|p-x\|^2 - \|r-x\|^2 - \omega^2 \cdot \|p-r\|^2. \end{aligned}$$

Note that $\angle prx > \pi/2$ since B', and so r, is contained in the diametral ball of px (because its center q' satisfies $\|q'-y\| \le (1/4)(c_1\varepsilon f(p)/2)$, and its radius is $\|p-q'\| \cdot \sin(\pi/20) \le (1/4)(c_1\varepsilon f(p)/2)$). Thus, $\|p-x\|^2 - \|r-x\|^2 > \|p-r\|^2$, which implies that $\pi_p(x) - \pi_r(x) > 0$. But then $x \notin V_p$, contradicting the assumption that $x \in C_p \cap V_p - B$. ◻

A.3 Proof of Lemma 4 *Proof.* The edge pq is a C_p-edge for some $\widehat{S}$. So V_{pq} intersects C_p. (Note that we do not know whether V_{pq} intersects C_q.) Take any point $x \in C_p \cap V_{pq}$. We claim that $\|q-x\| \le \|p-x\|/\sqrt{1-4\omega^2}$.

If $\|q-x\| \le \|p-x\|$, we are done. Suppose not. Then $\|p-q\| \le 2\,\|q-x\|$. Since $x \in V_{pq}$, we have $\|p-x\|^2 - P^2 = \|q-x\|^2 - Q^2$. After rearranginng terms, we get

$$\begin{aligned} \|p-x\|^2 &\ge \|q-x\|^2 - Q^2 \\ &\ge \|q-x\|^2 - \omega^2\,\|p-q\|^2 \\ &\ge \|q-x\|^2 - 4\omega^2\,\|q-x\|^2. \end{aligned}$$

This proves that our claim.

By Lemma 3, $\|p-x\| \le c_1\varepsilon f(p)$. Thus, by triangle inequality, $\|p-q\| \le \|p-x\| + \|q-x\| \le c_1(1+1/\sqrt{1-4\omega^2})\varepsilon f(p)$. ◻

A.4 Proof of Lemma 5 *Proof.* By Lemma 4, any edge in G_p has length at most $c_2\varepsilon f(p)$. By the (ε,δ)-sampling, any edge in G_p has length at least $\delta f(p)$. Thus $\|p-q\| \le \frac{c_2\varepsilon}{\delta} \cdot \|p-r\|$. (Recall that ε/δ is a constant.) By Lemma 4 and Lemma 1(ii), $\angle(pq, \mathcal{T}_p) \le \arcsin(c_2\varepsilon/2)$, which is at most $c_2\varepsilon$ for sufficiently small ε. ◻

A.5 Proof of Lemma 7 *Proof.* Since τ is a C_p-simplex, V_τ intersects C_p. There is a ball B with center $z \in V_\tau \cap C_p$ and radius Z such that B is orthogonal to the vertices of

τ, and B is further than orthogonal from other weighted vertices in $\widehat{S}$. Observe that $R'_\tau \leq Z$. Thus it suffices to show that $Z \leq \rho L_\tau$.

We prove the lemma for the constant $\rho = 5\nu\varepsilon/\delta$. Assume to the contrary that $Z > 5\nu\varepsilon L_\tau/\delta$. By Lemma 5(i) and the (ε,δ)-sampling, $\nu L_\tau \geq \delta f(p)$. Thus $Z > 5\nu\varepsilon L_\tau/\delta \geq 5\varepsilon f(p)$. Let y be the point on pz such that $\|y-z\| = Z - 5\varepsilon f(p)$. Thus, the ball B' with center at y and of radius $5\varepsilon f(p)$ is contained in B. By Lemma 4, the edges of τ incident to p have length at most $c_2\varepsilon f(p)$. The weight property $[\omega]$ implies that the weight P of p is at most $\omega c_2 \varepsilon f(p)$. We conclude that

$$5\varepsilon f(p) < \|p-y\| \leq (5+\omega c_2)\varepsilon f(p)$$

(the upper bound follows because on py, B' and the weighted ball at p overlap). Let q be the orthogonal projection of y onto $\mathcal{T}_p$ and let q' be the point on $\mathcal{M}$ closest to q. By Lemma 2(ii), we have

$$\|q'-y\| \leq \frac{(5+\omega c_2)\varepsilon}{4} f(p).$$

Setting $\omega \leq 3/c_2$, then $\|p-y\| \leq 8\varepsilon f(p)$ and $\|q'-y\| \leq 2\varepsilon f(p)$. By Lipschitz property $f(q') \leq f(p)+\|p-q'\| \leq f(p)+\|p-y\|+\|y-q'\| \leq (1+10\varepsilon)f(p)$, and so $f(q') \leq 2f(p)$ for sufficiently small ε. It follows that we can place a ball B'' strictly inside B' with center q' and radius $\varepsilon f(q') \leq 2\varepsilon f(p)$ (since the radius of B' is $5\varepsilon f(p)$, its center is y and $\|q'-y\| \leq 2\varepsilon f(p)$). So B'' is inside B and B'' is empty as B is empty. But the (ε,δ)-sampling implies that B'' contains a sample, a contradiction. ◻

A.6 Proof of Lemma 8

Proof. Let o and z be the orthocenters of τ_p and τ, respectively. Let R'_{τ_p} be the orthoradius of τ_p. The distance between z and $\mathrm{aff}(\tau_p)$ is equal to $\|o-z\|$. Observe that z is the closest point to p in $\mathrm{aff}(V_\tau)$. Therefore, Lemma 3 implies that

$$\|p-z\| \leq c_1\varepsilon f(p).$$

By applying Lemma 7 to τ_p, we get $R'_{\tau_p} \leq \rho L_{\tau_p}$. Let q be a vertex of τ_p. Using Lemma 5(i) and Lemma 4, we get

$$R'_{\tau_p} \leq \rho\nu \cdot \|p-q\| \leq c_2\rho\nu\varepsilon f(p).$$

Then triangle inequality implies that

$$\begin{aligned}
&\|o-z\| \\
&\leq \|p-z\| + \|p-q\| + \|q-o\| \\
&\leq c_1\varepsilon f(p) + c_2\varepsilon f(p) + \sqrt{Q^2 + R'^2_{\tau'}} \\
&\leq c_1\varepsilon f(p) + c_2\varepsilon f(p) + \sqrt{\omega^2\|p-q\|^2 + c_2^2\rho^2\nu^2\varepsilon^2 f(p)^2} \\
&\leq c_1\varepsilon f(p) + c_2\varepsilon f(p) + \sqrt{c_2^2\omega^2\varepsilon^2 f(p)^2 + c_2^2\rho^2\nu^2\varepsilon^2 f(p)^2},
\end{aligned}$$

which is at most $c_3\varepsilon f(p)$ for $c_3 = c_1 + c_2(1+\omega+\rho\nu)$. ◻

A.7 Proof of Lemma 9

Proof. First, $\mathrm{vol}(\tau) = \mathrm{vol}(\tau_q)\cdot D_q/j$. If $\mathrm{vol}(\tau) < \sigma L_\tau \mathrm{vol}(\tau_q)$ then $D_q < j\sigma L_\tau$. If $\mathrm{vol}(\tau) \geq \sigma L_\tau \mathrm{vol}(\tau_q)$ then $D_q \geq j\sigma L_\tau$. ◻

A.8 Proof of Lemma 18

Let k be the dimension of $\mathcal{M}$. Let σ be a j-cell of $\mathrm{Vor}_{\mathcal{M}}\ \hat{S}$ which is the intersection of the $(d-k+j)$-cell σ' of Vor $\hat{S}$ with $\mathcal{M}$. Let x be in the interior of σ, $L' = \mathrm{aff}(\sigma')$ and $\mathcal{N}'_x$ be the projection of $\mathcal{N}_x$ onto $L'-x$ (translation of L' so that x coincides with the origin). Let p be a sample point determining σ –that is, a vertex of the dual simplex σ^*–, because of the normal approximation of simplices, for any normal $\vec{n}_p$ of p there is a normal $\vec{n}_{\sigma^*}$ to σ^* with $\angle\vec{n}_p\vec{n}_{\sigma^*} \leq \pi/32$. This means that $\angle\mathcal{N}'_x\mathcal{N}_p \leq \pi/32$. With sufficiently small ε, normal variation on $\mathcal{M}$ implies that $\angle\mathcal{N}_p\mathcal{N}_x \leq \pi/32$. Therefore, $\angle\mathcal{N}_x\mathcal{N}'_x \leq \angle\mathcal{N}_x\mathcal{N}_p + \angle\mathcal{N}_p\mathcal{N}'_x \leq \pi/16$ holds.

For x in the interior of a restricted Voronoi cell σ, let $B(x)$ be a ball centered at x and with the smallest radius $r(x)$ so that $B(x)$ contains σ. We have $r(x) \leq c_5\varepsilon f(x)$. For a point $x \in \mathcal{M}$, let $\mathcal{B}(x)$ be interior of the union of all balls of radius $R(x) = f(x)$ that are tangent to $\mathcal{M}$ at x. Note that the intersection of $\mathcal{B}(x)$ and $\mathcal{M}$ is empty.

Proof. (*Sketch*) We use the notation established in the paragraphs above. The proof is by induction on j. For $j=0$, let σ' be a $(d-k)$-cell and suppose it intersects $\mathcal{M}$ in at least a point x. In this case $\mathcal{N}'_x = \mathrm{aff}(\sigma')$. Since $\sigma' \cap \mathcal{M}$ must lie in $B(x)$, and $\angle\mathcal{N}_x\mathcal{N}'_x \leq \pi/16$, it follows that, except x, σ' lies inside $\mathcal{B}(x)$ and so x is the unique intersection between σ' and $\mathcal{M}$.

Now, for $j \geq 1$, consider a j-cell σ and its corresponding $(d-k+j)$-cell σ', let $L' = \mathrm{aff}(\sigma')$, $\mathcal{M}'$ be the intersection of $\mathcal{M}$ and L', x be an arbitrary point in the interior of σ, $\mathcal{N}'_x$ be the orthogonal projection of $\mathcal{N}_x$ onto $L'-x$, and $\mathcal{T}'_x$ be the intersection of $\mathcal{T}_x$ with $L'-x$. $\mathcal{N}'_x$ and $\mathcal{T}'_x$ are the normal and tangent spaces of $\mathcal{M}'$ at x in L'. As noted above, $\angle\mathcal{N}_x\mathcal{N}'_x \leq \pi/16$.

Let $\vec{t}$ be a unit vector in $\mathcal{T}'_x$ (a direction) and let $F = F_{x,\vec{t}}$ be the $(d-k+1)$-half-flat determined by $\mathcal{N}'_x$ and positive factors of $\vec{t}$. Because $\angle\mathcal{N}_x\mathcal{N}'_x \leq \pi/16$, $\mathcal{M}_F = F \cap \mathcal{M}' \cap B(x)$ is a smooth 1-manifold in $B(x)$ (the intersection is transversal and hence generic inside $B(x)$) with boundary only possible on the boundary of $B(x)$ and on $x + \mathcal{N}'_x$. The normal space $\mathcal{N}''_y$ at an interior point $y \in \mathcal{M}_F$ in F is close to $\mathcal{N}_y$: for $\vec{n}''_y \in \mathcal{N}''_y$, there is an $\vec{n}_y \in \mathcal{N}_y$ with $\angle \vec{n}''_y \vec{n}_y \leq \pi/8$ (including the error between $\mathcal{N}_y$ and $\mathcal{N}'_y$ and between $\mathcal{N}'_y$ and $\mathcal{N}'_x$ –which determines F). The tangent $\vec{t}''_y$ on the other hand is close to $\vec{t}$: $\angle \vec{t}''_y \vec{t} \leq \pi/8$. $\mathcal{M}_F$ intersects the facets of $\sigma' \cap F$ almost orthogonally.

We now verify that it is possible to parametrize $\mathcal{M}_F$ with the distance from x; that is, that for $0 < r \leq r(x)$, there is a unique point $\gamma_{\vec{t}}(r)$ on $\mathcal{M}_F$ at distance r from x. Because $\mathcal{M}_F$ is a smooth 1-manifold, for r sufficiently small $\mathcal{M}_F$ is monotone with respect to r (that is, for each r, there is a unique y on $\mathcal{M}_F$ at distance r from x). We claim that $\mathcal{M}_F$ is monotone with respect to r inside $B(x)$. Otherwise some $y \in \mathcal{M}_F$ would have a normal $\vec{n}''_y$ in the direction of $y - x$, that is $\angle \vec{n}''_y (y - x) = 0$, which is a contradiction: there is a normal $\vec{n}_y$ in $\mathcal{N}_y$ such that $\angle \vec{n}_y \vec{n}''_y \leq \pi/8$, while $\angle \vec{t}(y - x) \leq \pi/8$ and $\angle \vec{t}\vec{n}_y \geq \pi/2 - \pi/8$, so this is a contradiction (making ε sufficiently small).

So let $\gamma_{\vec{t}}(r)$, $r \in I = [0, r(x)]$ be the parametrization of $\mathcal{M}_F$ according to its distance r from x. We claim that

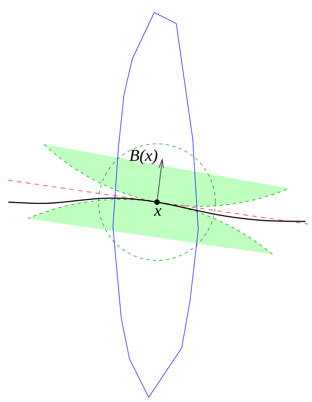

Figure 6: Within $B(x)$, $\mathcal{M}$ is "sandwiched" between large tangent empty balls.

$\gamma_{\vec{t}}$ starts at x inside $\sigma' \cap F$, eventually leaves $\sigma' \cap F$, and does not reenter it. It leaves $\sigma' \cap F$ because $\mathcal{M}_F$ may have a boundary only on the boundary of $B(x)$. To see that it does not reenter $\sigma' \cap F$, let $r^*_{\vec{t}}$ be such that $y^*_{\vec{t}} = \gamma_{\vec{t}}(r^*_{\vec{t}})$ is the first point of $\gamma_{\vec{t}}$ on the boundary of $\sigma' \cap F$, and let ω be the facet of σ' in Vor $\hat{S}$ intersected at this point. As observed above, $\gamma_{\vec{t}}$ intersects $\omega \cap F$ within $\pi/16$ from orthogonality and $\angle \vec{t}''_y \vec{t} \leq \pi/16$, so after $y^*_{\vec{t}}$, $\gamma_{\vec{t}}$ lies inside a cone with apex $y^*_{\vec{t}}$, axis parallel to the tangent to $\gamma_{\vec{t}}$ at $y^*_{\vec{t}}$, and aperture $\pi/16$. It follows by convexity of $\sigma' \cap F$ that, after $y^*_{\vec{t}}$, $\gamma_{\vec{t}}$ is outside of $\sigma' \cap F$, that is, it does not reenter σ'.

We claim that $h(\vec{t}) = y^*_{\vec{t}}$ defines a homeomorphism between the $(j-1)$-sphere and the intersection of $\mathcal{M}'$ with the boundary of σ'. The map h is continuous because the intersection of $\mathcal{M}$ with each $(d-k+j-1)$-dimensional cell on the boundary of σ' is a topological ball (by inductive argument), and these pieces are glued continuously. Since h is one-to-one and continuos, it is a homeomorphism (as it is a map between compact spaces).

Thus, we can then use the fibers $\gamma^*_{\vec{t}} = \gamma_{\vec{t}}([0, r^*_{\vec{t}}])$ to define a homeomorphism between a (Euclidean) ball and σ in the natural way (proceeding as in the standard proof that a convex cell is a topological ball): a ray of the ball in the direction $\vec{t}$ maps to the fiber $\gamma^*_{\vec{t}}$ in the direction $\vec{t}$. ◻