HKUST Theoretical Computer Science Center Research Report HKUST-TCSC-2004-05

Surface Reconstruction from Noisy Samples

Siu-Wing Cheng* Sheung-Hung Poon*

August 3, 2004

Abstract

We study the problem of reconstructing a closed smooth surface S in R^3 from n noisy samples. We present a *deterministic* model of noisy samples. We justify it by proving that the model is satisfied with high probability, if we randomly perturb dense uniform samples on S in the normal directions. Let $0 < \gamma \leq 1/8$ be a constant parameter chosen in advance. We develop an $O(n^{2+\gamma})$-time algorithm to construct a triangulated surface based on the deterministic model. As $n \to \infty$, the output surface is homeomorphic to S and its approximation error converges to zero.

1 Introduction

The problems of reconstructing surfaces from sample points has been studied extensively in computer graphics, geometric modeling, image processing, computer vision, and computational geometry. The input consists of sample points from the unknown smooth closed surface S. The problem calls for computing a polygonal surface that is provably *faithful*. That means, when the sampling density is sufficiently large, the reconstructed polygonal surface should be homeomorphic to S and the approximation error should approach zero. There are two types of approximation error. First, the Hausdorff distance between the reconstruction and S. Second, the differences between the normals of the reconstruction and that of S. Both types of error should approach zero as the sampling density increases.

The first and widely known reconstruction algorithm in computer graphics is the work of Hoppe et al. [18, 20, 19]. Their work performs quite well in practice, and allows the presence of noise. Later, Curless and Levoy [8] presented a volumetric method to reconstruct surface from range images. Medioni and Tang [21, 24] applied a technique called tensor voting to infer the surface and features on it. All these algorithms can also handle noise in the input, and appear to be quite successful in practice. However, they do not have theoretical guarantees.

Some faithful reconstruction algorithms were proposed in computational geometry for the simple case that sample points are noiseless. Edelsbrunner and Mucke [12] used the α-shape to construct the surface, and Edelsbrunner [11] proposed another surface reconstruction algorithm by wrapping the sample points. The main drawback is that they perform well only for uniform sample set. Amenta and Bern [1, 2] extended the approach of Voronoi filtering to do surface reconstruction. Later, Amenta,

*Department of Computer Science, HKUST, Hong Kong. Email:{scheng,hung}@cs.ust.hk. Research partly supported by Research Grant Council, Hong Kong, China (project no. HKUST 6169/03E).

Choi, Dey, and Leekha [3] presented a simpler algorithm with a proof that the reconstruction is homeomorphic to the original surface. Amenta et al. [4, 5] gave another faithful reconstruction algorithm, which is called power crust. One merit of the power crust algorithm is that it always output a water-tight surface regardless of the sampling density of the input data set. Another algorithm given by Boissonnat and Cazals [6] reconstructs a smooth implicit surface interpolating the sample points. One can then mesh the implicit surface to obtain a mesh of the desired resolution.

As noise often arises in practical sample point set, faithful surface reconstruction algorithms are needed for sample set in the presence of noise. Recently, Mitra and Nguyen [23] presented methods to estimate surface normals from point clouds, but did not proceed to finally reconstruct the surface for the sample points. Cheng et al. [7] proposed a probabilistic noise model for curve reconstruction in two dimensions. Based on this model, they developed a reconstruction algorithm the output of which is faithful with high probability. This is the first theoretical result known for the curve reconstruction problem in the presence of noise. In this paper, we extend our approach in Cheng et al. [7] to do surface reconstruction from noisy samples with theoretical guarantees in three dimensions. Not only do we improve the algorithm to make it run faster, we also make the noise model deterministic which extends its applicability and simplifies the analysis of the algorithm. Very recently and independently, Dey and Goswami [10] proposed a surface reconstruction algorithm from noisy samples. There are several major differences between their work and ours in this thesis. In their noise model, the noise amplitude is proportional to the local feature size and inversely proportional to the sampling density; whereas in our model, it is an independent constant. Their algorithm selects a set of "boundary" noisy sample points for reconstruction; whereas our algorithm tries to estimate some center points close to the original surface for reconstruction. Finally, they gave an implementation for their algorithm and had some promising results, but we havn't implemented our algorithm.

We gave an introduction for the and surface reconstruction problem in the current section. We introduce some basic notations and some basic geometric lemmas in Section 2. In Section 3, we present a deterministic noise model for the distribution of sample points on surfaces and also justify it. Our surface reconstruction algorithm is described in section 4. Then we give an overview on the analysis of our faithfulness proof in Section 5, and we give the details of the analysis in Section 6, 7 and 8. We summarize our result in Section 9. Note that the proofs of some technical lemmas are given in the appendix.

2 Preliminaries

We use S to denote the smooth closed surface and P to denote the set of noisy sample points from S. The *noise amplitude* of P is the maximum of the distances from P to the surface S. We denote the noise amplitude of P to S by δ. We suppose $\delta < 1$ since we only consider noise is close to the original surface S.

The *medial axis* of S is the set of centers of empty balls that touch S at more than one point, and for any pint $x \in S$, the *local feature size* $f(x)$ is the minimum distance from x to the medial axis of S. We assume that $\min_{x \in S} f(x) = 1$ for convenience. So for any $x \in S$, $f(x) \geq 1$.

For any $\alpha \in [-\delta, \delta]$, S_α^+ (resp., S_α^-) is the collection of points in R^3 that lie outside (resp., inside)

S and their distances from S are exactly α. We use S_α when it is unimportant to distinguish between inside and outside. Note that S_α and S have the same medial axis. We use **S** to denote the set of points sandwiched between S_δ^+ and S_δ^-.

Given two objects A and B, we use $d(A, B)$ to denote the minimum distance between A and B. Given two coplanar vectors/lines in 3D, the plane *spanned by them* is the plane that contains them. Given two vectors, lines, or planes A and B in 3D, we use $\angle(A, B)$ to denote the smallest angle between them. If A and B are disjoint, we translate one of them to measure $\angle(A, B)$.

For each point $p \in S_\alpha$, we use n_p to denote the *outward unit normal* of S_α at p, and T_p to denote the *tangent plane* of S_α at p. We use $\tilde{p}$ to denote the point on S such that $d(p, \tilde{p}) = \alpha$. Note that n_p and $n_{\tilde{p}}$ have the same support line, and T_p and $T_{\tilde{p}}$ are parallel. We use $B(p, d)$ to denote a ball centered at p with radius d.

We state three geometric lemmas that will be useful later. Their proofs can be found in the appendix.

Lemma 2.1 *Any point p on S_α has two tangent balls with radii $f(\tilde{p}) - \alpha$ whose interior do not intersect S_α.*

The next lemma follows from results in [1, 15]. It basically says that the neighborhood of any point on S_α is fairly flat, and the normal varies slowly.

Lemma 2.2 *Assume that $\delta \leq 1/6$. Let $p, q \in S_\alpha$ be two points.*

(i) If $d(p, q) = \lambda(f(\tilde{p}) - \delta)$ for some $\lambda < 1$, then $d(q, T_p) \leq d(p, q)^2$ and $\angle(pq, T_p) \leq \arcsin(d(p, q))$.

(ii) If $d(p, q) = \lambda \min\{f(\tilde{p}) - \delta, f(\tilde{q}) - \delta\}$ for some $\lambda \leq 1/5$, then $\angle(n_p, n_q) \leq 3d(p, q)$.

The following lemma shows that the distance between nearby points on S are preserved in S_α up to a constant factor.

Lemma 2.3 *Assume $\delta \leq 1/10$. Let $p, q \in S_\alpha$ be two points such that $d(\tilde{p}, \tilde{q}) = d$ for some $d \leq 1/5$. Then $d/2 \leq d(p, q) \leq 3d/2$.*

3 Sampling and noise model

3.1 Deterministic noise model

Let ω be a fixed positive constant (we think of it as very small). We assume that the set P of noisy samples satisfies the following conditions.

There exists parameters $0 < \epsilon_2 \leq \epsilon \leq \epsilon_1 < 1$ depending on n such that

(i) $\epsilon_1 \to 0$ as $n \to \infty$,

(ii) $\epsilon_1 = O(\epsilon)$, $\epsilon_2 = \Omega(\epsilon^c)$ for some constant $c \geq 1$.

and for any point $p \in$ **S**, we have

(i) $|B(p,\epsilon) \cap P| \geq 1$,

(ii) $|B(p,\epsilon_1) \cap P| \geq 2\ln^\omega n$,

(iii) $|B(p,\epsilon_2) \cap P| \leq \frac{1}{2}\ln^\omega n$.

In Theorem 3.1, specifically, we can prove that for $\epsilon = \Theta(\sqrt[3]{\frac{\ln^\omega n}{n}})$, $\epsilon_1 = \Theta(\sqrt[3]{\frac{\ln^\omega n}{n}})$ and $\epsilon_2 = \Theta(\sqrt{\frac{\ln^\omega n}{n}})$, our deterministic model is satisfied with high probability by a set of sample points generated according to a uniform distribution. In a sense, the parameters $\epsilon, \epsilon_1, \epsilon_2$ measure the "sampling density". It would be nice that the three parameters collapse to a single one (with some appropriate adjustment in the conditions), but our analysis in the justification is not strong enough to do so. It is also more natural that $\epsilon_2 = \Omega(\epsilon)$, but we can only justify our model for $c \geq 3/2$. Therefore, we will assume throughout this paper that $\epsilon_2 = \Omega(\epsilon^{3/2})$. We remark that our algorithm does not know the values of ϵ, ϵ_1, ϵ_2, and the noise amplitude δ.

3.2 Justification

Consider generating the noisy samples using the following random process. First, points are drawn from S according to a uniform distribution. That is, the probability of drawing a point from a region A is equal to $\frac{\mathrm{area}(A)}{\mathrm{area}(S)}$. Second, each point x drawn is perturbed uniformly within a line segment with width 2δ, centered at x, and aligned with n_x. Moreover, the distribution of each sample point is independently identical. A similar probabilistic model was used for the curve case in [7]. Our goal is to prove that the samples generated by the above random process satisfy our deterministic model with probability at least $1 - \exp(-\Omega(\ln^\omega n))$.

For any constant $k > 0$, let λ_k denote $\sqrt[3]{\frac{k\ln^\omega n}{n}}$. Note that $0 < \lambda_k < 1$ for sufficiently large n. Let s be a point on S_α. Let $R_1 \subseteq S_\alpha$ denote the patch $\{p \in S_\alpha : d(p,s) \leq \lambda_k\}$. Let $R_2 \subseteq S_\alpha$ denote the patch $\{p \in S_\alpha : d(p,s) \leq \lambda_k^{3/2}\}$. The λ_k*-cell at* s is the union of line segments with midpoints on R_1, normal to R_1, and with length $2\lambda_k\delta$. Note that the λ_k-cell may not lie completely inside **S**, but at least half of it on one side of R_1 does. The λ_k*-rod at* s is the union of line segments with length 2δ, inside **S**, and normal to some point in R_2. The following two lemmas bound the sizes of a λ_k-cell and a λ-rod.

Lemma 3.1 *The* λ_k*-cell at a point* s *lies inside* $B(s, 2\lambda_k)$.

Proof. Assume that $s \in S_\alpha$. Let C denote the λ_k-cell at s. Let p be any point in C. Let q be the projection of p along the normal to S_α, i.e., $\tilde{p} = \tilde{q}$. Then $d(p,q) \leq \lambda_k\delta$ and $d(q,s) \leq \lambda_k$. Thus $d(p,s) \leq d(p,q) + d(q,s) \leq \lambda_k\delta + \lambda_k \leq 2\lambda_k$. ◻

Lemma 3.2 *The* λ_k*-rod at a point* s *contains* $B(s, \lambda_k^{3/2}/4) \cap \mathbf{S}$.

Proof. Let C denote the λ_k-rod at s. Let p be a point in $B(s, \lambda_k^{3/2}/4) \cap \mathbf{S}$. Assume that $p \in S_\alpha$. Let s' denote the projection of s onto S_α along the normal direction. Clearly $d(s', s) \leq d(s,p) \leq \lambda_k^{3/2}/4$. Thus $d(s',p) \leq d(s',s) + d(s,p) \leq \lambda_k^{3/2}/2$. Assume that $s \in S_\beta$. Let p' denote the projection of p onto S_β along the normal direction. By Lemma 2.3, $d(s,p') \leq 2d(s',p) \leq \lambda_k^{3/2}$. Thus $p' \in C \cap S_\beta$, which implies that $p \in C \cap S_\alpha$. ◻

We give a highlight of our proof. We first upper and lower bound the probabilities of a sample appearing in the λ_k-cell and λ_k-rod at s. Based on these probability bounds, we can show that the number of samples inside the λ_k-cell at s is essentially $\Theta(\ln^\omega n)$ with high probability. The same holds for the λ_k-rod at s. By setting ϵ_1 and ϵ_2 appropriately, we then show that the λ_k-cell lies inside the ball $B(s, \epsilon_1)$ and the ball $B(s, \epsilon_2)$ lies inside the λ_k-rod. Hence the conditions in our deterministic noise model are satisfied. We give the details of the proof below.

The next technical lemma follows from the fact that a small neighborhood of any point on S_α is fairly flat.

Lemma 3.3 *Assume that $r \leq (1-\delta)/20$. Let p be a point on S_α. Let Y be the infinite cylinder with radius r and axis aligned with n_p. Then the area of the connected component in $Y \cap S_\alpha$ containing p is $\Theta(r^2)$.*

Next, we bound the probability of a sample appearing in the λ_k-cell and λ_k-rod.

Lemma 3.4 *Let C be the λ_k-cell or λ_k-rod at a point $s \in \mathbf{S}$. Assume that n is so large that $\lambda_k \leq 1/4$. Then for any sample p, $\kappa_2\lambda_k^3 \leq \Pr(p \in C) \leq \kappa_1\lambda_k^3$ for some constants κ_1 and κ_2.*

Proof. Consider the case that C is a λ_k-cell. Assume that s lies on S_α. Let $R \subseteq S$ denote the patch $\{\tilde{x} : x \in C \cap S_\alpha\}$. Project R orthogonally onto the region R' on the tangent plane $T_{\tilde{s}}$ at $\tilde{s}$. Take a point q in the boundary of $C \cap S_\alpha$. Let q' denote the orthogonal projection of $\tilde{q}$ onto $T_{\tilde{s}}$. By definition, $d(s, q) = \lambda_k$. Then Lemma 2.3 implies that $\lambda_k/2 \leq d(\tilde{s}, \tilde{q}) \leq 2\lambda_k$.

Clearly, $d(\tilde{s}, q') \leq d(\tilde{s}, \tilde{q}) \leq 2\lambda_k$. Thus R' lies within the disk on $T_{\tilde{s}}$ centered at $\tilde{s}$ with radius $2\lambda_k$. Then by Lemma 3.3, $\mathrm{area}(R) = O(\lambda_k^2)$. On the other hand, $d(\tilde{s}, q') = d(\tilde{s}, \tilde{q}) \cos\angle(\tilde{q}\tilde{s}, T_{\tilde{s}}) \geq \lambda_k/2 \cdot \cos\angle(\tilde{q}\tilde{s}, T_{\tilde{s}})$. By Lemma 2.2(i), $\angle(\tilde{q}\tilde{s}, T_{\tilde{s}}) \leq \arcsin(d(\tilde{s}, \tilde{q})) \leq 2d(\tilde{s}, \tilde{q}) \leq 4\lambda_k$. Since $\lambda_k \leq 1/4$ by assumption, $\angle(\tilde{q}\tilde{s}, T_{\tilde{s}}) \leq 1$. So $\cos\angle(\tilde{q}\tilde{s}, T_{\tilde{s}}) \geq \cos(1) \geq 1/2$. This implies that $d(\tilde{s}, q') \geq \lambda_k/4$. Thus R' contains a disk on $T_{\tilde{s}}$ centered at $\tilde{s}$ with radius $\lambda_k/4$. Combining with the previous upper bound, we conclude that $\mathrm{area}(R) = \Theta(\lambda_k^2)$.

The probability of drawing a point from R is thus $\Theta(\lambda_k^2)/\mathrm{area}(S)$. The probability that the perturbation throws the point drawn from R into C is at least $\frac{\lambda_k\delta}{2\delta} = \frac{\lambda_k}{2}$, and at most $\frac{2\lambda_k\delta}{2\delta} = \lambda_k$. It follows that for any sample p, $\Pr(p \in C) = \Theta(\lambda_k^3)$, assuming that $\mathrm{area}(S)$ is an intrinsic constant for S.

Now consider the case that C is a λ_k-rod. We also let $R \subseteq S$ denote the patch $\{\tilde{x} : x \in C \cap S_\alpha\}$. By the definition of a λ_k-rod and an argument similar to the one in the above, we can show that $\mathrm{area}(R) = \Theta(\lambda_k^3)$. Then the probability of a sample falling inside C is same as the probability of drawing a point from R, which is $\mathrm{area}(R)/\mathrm{area}(S) = \Theta(\lambda_k^3)$. ◻

Next, we bound the number of samples inside the λ_k-cell and λ_k-rod in Lemma 3.6. The following Chernoff bound [16] stated in Lemma 3.5 will be needed for doing that.

Lemma 3.5 *Let the random variables $X_1, X_2, \ldots, X_n$ be independent, with $0 \leq X_i \leq 1$ for each i. Let $S_n = \sum_{i=1}^n X_i$, and let $E(S_n)$ be the expected value of S_n. Then for any $\sigma > 0$, $\Pr(S_n \leq (1-\sigma)E(S_n)) \leq \exp(-\frac{\sigma^2 E(S_n)}{2})$, and $\Pr(S_n \geq (1+\sigma)E(S_n)) \leq \exp(-\frac{\sigma^2 E(S_n)}{2(1+\sigma/3)})$.*

Lemma 3.6 *Assume that n is so large that $\lambda_k \leq 1/4$. Let C be a λ_k-rod or λ_k-cell at a point $s \in$* **S**. *Let κ_1 and κ_2 be the constants in Lemma 3.4.*

(i) C is non-empty with probability at least $1 - \exp(-\Omega(\ln^{\omega} n))$.

(ii) For any constant $\kappa > \kappa_1 k$, the number of samples in C is at most $\kappa \ln^{\omega} n$ with probability at least $1 - \exp(-\Omega(\ln^{\omega} n))$.

(iii) For any constant $\kappa < \kappa_2 k$, the number of samples in C is at least $\kappa \ln^{\omega} n$ with probability at least $1 - \exp(-\Omega(\ln^{\omega} n))$.

Proof. Let $X_i (i = 1, \ldots, n)$ be a random binomial variable taking value 1 if the sample point s_i is inside C, and value 0 otherwise. Let $S_n = \sum_{i=1}^{n} X_i$. Then $E(S_n) = \sum_{i=1}^{n} E(X_i) = n \cdot \Pr(s_i \in C)$. This implies that

$$E(S_n) \leq \kappa_1 n \lambda_k^3 = \kappa_1 k \ln^{\omega} n, \text{ and}$$
$$E(S_n) \geq \kappa_2 n \lambda_k^3 = \kappa_2 k \ln^{\omega} n$$

By Lemma 3.5,

$$\begin{aligned} \Pr(S_n \leq 0) &= \Pr(S_n \leq (1-1)E(S_n)) \\ &\leq \exp(-\frac{E(S_n)}{2}) \\ &\leq \exp(-\Omega(\ln^{\omega} n)). \end{aligned}$$

Consider (ii). Let $\sigma = \frac{\kappa}{\kappa_1 k} - 1 > 0$. We have

$$\kappa \ln^{\omega} n = \kappa_1 n \lambda_k^3 (1+\sigma) \geq (1+\sigma) E(S_n).$$

By Lemma 3.5,

$$\begin{aligned} \Pr(S_n > \kappa \ln^{\omega} n) &\leq \Pr(S_n > (1+\sigma)E(S_n)) \\ &\leq \exp(-\frac{\sigma^2 E(S_n)}{2 + 2\sigma/3}) \\ &= \exp(-\Omega(\ln^{\omega} n)). \end{aligned}$$

Consider (iii). Let $\sigma = 1 - \frac{\kappa}{\kappa_2 k} > 0$. We have

$$\kappa \ln^{\omega} n = \kappa_2 n \lambda_k^3 (1-\sigma) \leq (1-\sigma) E(S_n).$$

By Lemma 3.5,

$$\begin{aligned} \Pr(S_n < \kappa \ln^{\omega} n) &\leq \Pr(S_n < (1-\sigma)E(S_n)) \\ &\leq \exp(-\frac{\sigma^2 E(S_n)}{2}) \\ &= \exp(-\Omega(\ln^{\omega} n)). \end{aligned}$$

□

We are ready to prove that the random process produces samples that satisfy our deterministic model with high probability.

Theorem 3.1 *Let X be n points on S that are drawn uniformly. For each $x \in X$, define $p_x = x + \alpha_x n_x$, where α_x is drawn uniformly and independently from $[-\delta, \delta]$. There exists $\epsilon, \epsilon_1, \epsilon_2$ such that $\{p_x : x \in X\}$ satisfy our deterministic model with probability at least $1 - \exp(-\Omega(\ln^\omega n))$.*

Proof. Choose two constants a and b such that $2 < \kappa_2 a$ and $\frac{1}{2} > \kappa_1 b$. We prove the theorem for $\epsilon = 2\lambda_1$, $\epsilon_1 = 2\lambda_a$, and $\epsilon_2 = \lambda_b^{3/2}/4$. First of all, it is clear that $\epsilon_1 = O(\epsilon)$ and $\epsilon_2 = \Omega(\epsilon^{3/2})$. For convenience, we use "with high probability" to mean "with probability at least $1 - \exp(-\Omega(\ln^\omega n))$".

Let p be a point in **S**. Let C be the λ_1-cell at p. By Lemma 3.6(i), C contains a sample with high probability. By Lemma 3.1, C lies inside the ball $B(p, 2\lambda_1) = B(p, \epsilon)$. Thus $|B(p, \epsilon) \cap P| \geq 1$ with high probability.

Let C_1 be the λ_a-cell at p. By Lemma 3.6(iii), C_1 contains at least $2\ln^\omega n$ samples with high probability. By Lemma 3.1, C_1 lies inside the ball $B(p, 2\lambda_a) = B(p, \epsilon_1)$. Thus $|B(p, \epsilon_1) \cap P| \geq 2\ln^\omega n$ with high probability.

Let C_2 be the λ_b-rod at p. By Lemma 3.6(ii), C_2 contains at most $\frac{1}{2}\ln^\omega n$ samples with high probability. By Lemma 3.2, the ball $B(p, \lambda_b^{3/2}/4) \cap \mathbf{S} = B(p, \epsilon_2) \cap \mathbf{S}$ lies inside C_2. Thus $|B(p, \epsilon_2) \cap P| \leq \frac{1}{2}\ln^\omega n$ with high probability. $\square$

4 Algorithm

From now on, we suppose the input sample points satisfy the deterministic Similar to the framework of the curve reconstruction algorithm in Cheng et al. [7], our algorithm for surface consists of three main steps called POINT ESTIMATION, PRUNING and SURFACE RECONSTRUCTION. Our algorithm first filters noise and then applies cocone to obtain a triangulated surface. The noise filtering procedure consists of two steps, POINT ESTIMATION and PRUNING. In the first step, which is the POINT ESTIMATION step, new points are computed that are provably close to S. However, the distances of these new points from S can be much larger than the distances among them. Using all the new points to form a reconstruction would produce a highly rugged surface. In the second step, which is the PRUNING step, some of the new points are pruned so that the interpoint distances in the remaining subset of new points are large relative to their distances from S. Now we can apply the cocone algorithm (or any faithful surface reconstruction algorithm for noiseless samples) to produce a provably good surface in the final SURFACE RECONSTRUCTION step. We remark that our algorithm does not estimate the model parameters $\epsilon, \epsilon_1, \epsilon_2$. It may be possible to do so, but we do not rely on any such estimate.

We further highlight the POINT ESTIMATION step. At each sample $p \in P$, we grow a ball neighborhood called *coarse neighborhood* until the noisy samples inside the neighborhood fit inside a sufficiently thin slab. This thin slab is an approximation of the tangent plane at $\tilde{p}$, i.e., the direction normal to the thin slab approximates n_p. However, due to noise, the size of the ball neighborhood is in the order of the noise amplitude δ. As the surface can bend quite a lot inside this ball neighborhood, the approximation error of the estimated normal is also in the order of δ. Thus we have a second phase to improve the estimate in a smaller neighborhood called *refined neighborhood* around p. We take some directions slightly different from the normal estimate due to the coarse neighborhood. We construct a narrow tube oriented in each direction. We trim each tube to the shortest cylinder containing all samples inside. Among all cylinders, the orientation of the shortest one provably approximates n_p well. We provide the algorithmic details below. Recall that $\omega > 0$, $\rho \geq 20$, and $0 < \gamma \leq 1/8$ are constants chosen in advance.

POINT ESTIMATION. For each noisy sample $p \in P$, we compute a point p^* that approximates $\tilde{p}$ as follows.

INITIAL BALL. We compute a ball $\mathit{initial}(p)$ centered at p with $\ln^\omega n$ samples inside.

COARSE NEIGHBORHOOD. We initialize $\mathit{coarse}(p)$ to be the ball centered at p with radius $r_p = \mathrm{radius}(\mathit{initial}(p))^\gamma$. We compute an infinite slab $\mathit{candidate}(p)$ that contains all samples inside $\mathit{coarse}(p)$. A slab of minimum width would be nice, but a 2-approximation suffices which can be computed in linear time [17]. We repeatedly double the size of $\mathit{coarse}(p)$ and maintain $\mathit{candidate}(p)$ until the radius of $\mathit{coarse}(p)$ is not less than $\rho \cdot \mathrm{width}(\mathit{candidate}(p))$.

REFINED NEIGHBORHOOD. Let $\mathcal{D}$ be the unit sphere of directions centered at p. We cover $\mathcal{D}$ by a set of caps as follows. Grid the boundary of the bounding box of $\mathcal{D}$ into square cells of width $\sqrt{r_p}$. Project the grid vertices onto $\mathcal{D}$ towards the center of $\mathcal{D}$ and put a cap with angular radius $\sqrt{r_p}$ centered at each projected grid vertex. Let h_p be the line through p orthogonal to $\mathit{candidate}(p)$. We examine all caps on $\mathcal{D}$. Let ℓ be the line through p and the center of a cap. If $\angle(h_p, \ell) > \pi/10$, examine the next cap. Otherwise, let $L_p(\ell)$ denote the infinite tube with axis ℓ and radius $r_p/4$. We trim $L_p(\ell)$ to obtain the shortest cylinder $C_p(\ell)$ that contains all noisy samples inside $L_p(\ell) \cap \mathit{coarse}(p)$. After going through all caps, we pick the shortest cylinder $C_p(\ell_p^*)$ generated. The line ℓ_p^* approximates n_p and the center p^* of $C_p(\ell_p^*)$ approximates $\tilde{p}$.

PRUNING. We scan the cylinder centers in an arbitrary order and select a subset C^*: when we select the center p^*, we delete all centers q^* not yet selected that satisfy $d(p^*, q^*) \leq \sqrt[8]{r_p/4}$.

SURFACE CONSTRUCTION. We run the cocone algorithm [3] on the centers in C^* and return the output surface, say N. One modification is that instead of using polar normals, we use the estimated normal ℓ_p^* for each $p^* \in C^*$. The details are provided in section 8.5.

We analyze the running time of the algorithm. PRUNING clearly takes $O(n^2)$ time. In SURFACE CONSTRUCTION, the most time-consuming step is the construction of the 3D Voronoi diagram which takes $O(n^2)$ time. INITIAL BALL takes $O(n)$ time for each sample. It remains to show that COARSE NEIGHBORHOOD and REFINED NEIGHBORHOOD take $O(n^{1+\gamma})$ time for each sample. We need a technical lemma on the relations among $n, \epsilon, \epsilon_1, \epsilon_2$, whose proof is given in the appendix. Recall that we are working with the deterministic noise model, so these relations are derived based on it.

Lemma 4.1 $\epsilon, \epsilon_1 = O((\frac{\ln^\omega n}{n})^{2/9})$ *and* $\epsilon_2 = \Omega((\frac{1}{n})^{3/4})$.

Lemma 4.2 *The algorithm runs in $O(n^{2+\gamma})$ time.*

Proof. In growing $coarse(p)$, we spend $O(n)$ time to compute $candidate(p)$ after each doubling. So the total time needed is $O(n\log(R_p/r_p))$, where R_p is the radius of the final $coarse(p)$. By the model, $\mathrm{radius}(initial(p)) \geq \epsilon_2$ which is $\Omega(1/n^{3/4})$ by Lemma 4.1. Thus, $r_p = \mathrm{radius}(initial(p))^\gamma = \Omega(1/n^{3\gamma/4})$. We will show in Lemma 6.2 that $R_p \leq 20\rho\delta+2r_p$, which is $O(1+r_p)$. Thus $\log(R_p/r_p) = O(\log n)$, implying that COARSE NEIGHBORHOOD takes $O(n\log n)$ time per sample. Thus it takes $O(n^2\log n)$ total time.

In REFINED NEIGHBORHOOD, for each $L_p(\ell)$, collecting the samples inside $L_p(\ell) \cap coarse(p)$ and computing $C_p(\ell)$ takes $O(n)$ time. Let K be the set of caps on $\mathcal{D}$ considered in the algorithm for sample point p. Since the angular radius of each cap is $\sqrt{r_p}$, the area of each cap is $O(r_p)$. Thus the number of caps $|K|$ considered for sample point p is $O(1/r_p)$. So the total time to compute all $C_p(\ell)$ for the corresponding directions in K is $O(n/r_p)$. As $r_p = \Omega(1/n^{3\gamma/4})$, we have $n/r_p = O(n^{1+3\gamma/4}) = O(n^{1+\gamma})$. Thus REFINED NEIGHBORHOOD takes $O(n^{1+\gamma})$ time per sample. Thus it takes $O(n^{2+\gamma})$ total time.

On the other hand, as PRUNING and SURFACE RECONSTRUCTION takes $O(n^2)$ time. Hence the total running time is $O(n^{2+\gamma})$. ◻

5 Overview of analysis

In the consequent sections, we will prove the theoretical guarantees ensured by our algorithm, that are summarized in the following theorem.

Theorem 5.1 *Assume $\rho > 1$ and $\delta \leq 1/(1600\rho^2)$. Let $0 < \gamma \leq 1/8$ be a constant parameter. Given a noisy sampling P of S that satisfies our deterministic noise model, our algorithm constructs a a triangulated surface N in $O(n^{2+\gamma})$ time such that as n increases,*

(i) each vertex p^* of N converges to $\tilde{p}$,

(ii) the normal of each triangle T in N converges to the normal at $\tilde{p}$ for some vertex p^* of T, and

(iii) N is homeomorphic to S.

Not surprisingly, the framework of our proofs for faithfulness for the reconstructed surface is similar to the curve case in [7]. We need to go through the analysis for the coarse neighborhood, the refined neighborhood, and the homeomorphism proof because the situation in three dimensions is much more complicated.

Consider the final coarse neighborhood $coarse(p)$ at a sample point p. Due to the criterion to sustain the growing of $coarse(p)$, $coarse(s)$ must have radius $\Theta(\rho\delta+r_p)$, where r_p is the radius of initial coarse ball at p. Next, we would like to argue that the normal of $candidate(s)$ approximates the surface normal at $\tilde{s}$. We claim that the deviation is actually $O(\rho\delta + r_p)$. We prove this by contradiction and assume that $candidate(s)$ is tilted a lot. Then a significant volume of $\mathbf{S} \cap coarse(p)$ lies outside $candidate(s)$. Our goal is to show that this particular volume contains a noisy sample so that a contradiction is established. The details are in Section 6.

As the normal estimation in the coarse neighborhood depends on the noise amplitude, and thus does not converge to zero when n increases. So we need to proceed to have a better estimate in the refined neighborhood. In the refined neighborhood at p, we rotate a thin cylinder of width $\Theta(r_p)$ centered at p. In the algorithm, the both bases of a cylinder are set to tightly bound all the sample points inside the cylinder. Our target cylinder axis direction ℓ_p^* is a direction so that the cylinder height is minimum. For efficiency, in our algorithm, we only seek for an approximate minimum height. The proof intuition is that when ℓ_p^* is close to the surface normal, the corresponding cylinder height will be quite small; otherwise the corresponding cylinder height will be large. Thus we can eliminate the latter option, and we can prove that ℓ_p^* is close to the surface normal. From this, we can then prove the center p^* of this cylinder of small height is close to the original surface S. The details are in Section 7.

Finally, to prove that the normals of the output triangles approximate well the surface normals and the reconstructed surface is homeomorphic to the original surface S, we adapts the proof technique in the cocone paper [3] for noiseless sample points. The main problem we need to get through is that in their original proof they assume that all the sample points must be on the original surface S, which is not true in our case. The details are in Section 8.

6 Coarse neighborhood

We first bound the size of $coarse(p)$ which allows us to bound the approximation error of the initial estimate of n_p.

6.1 Radii of $initial(p)$ and $coarse(p)$

The next two lemmas bound the radii of $initial(p)$ and $coarse(p)$.

Lemma 6.1 *Then $\Omega(\epsilon^{3/2}) = \epsilon_2 \leq \mathrm{radius}(initial(p)) \leq \epsilon_1 = O(\epsilon)$. Also, $r_p^2 \geq \epsilon$ and $r_p = O(\epsilon^\gamma)$ for sufficiently large n.*

Proof. By our deterministic model, as $initial(p)$ contains $\ln^\omega n$ samples, the radius of $initial(p)$ is less than $\epsilon_1 = O(\epsilon)$ and at least $\epsilon_2 = \Omega(\epsilon^{3/2})$. Then $r_p = (\mathrm{radius}(initial(p)))^\gamma = O(\epsilon^\gamma)$.

Since $r_p \geq \mathrm{radius}(initial(p))^\gamma \geq \epsilon_2^\gamma$ and $\epsilon_2 = \Omega(1/n^{3/4})$ by Lemma 4.1, we have $r_p = \Omega(1/n^{3\gamma/4}) = \Omega(1/n^{3/32})$ as $\gamma \leq 1/8$. By Lemma 4.1, $\epsilon = O((\frac{\ln^\omega n}{n})^{2/9})$. Thus for sufficiently large n, $r_p^2 \geq \epsilon$. ◻

Lemma 6.2 *Assume $\rho > 1$ and $\delta \leq 1/(1600\rho^2)$. Then*

$$\max\{2\sqrt{\rho}\delta, r_p\} \leq \mathrm{radius}(coarse(p)) \leq 10\rho\delta + 2r_p$$

for sufficiently large n.

Proof. Let $R_p = \mathrm{radius}(coarse(p))$. We first upper bound R_p. Suppose to the contrary that $R_p > 10\rho\delta + 2r_p$. Since the final $coarse(p)$ is obtained by repeated doubling $B(p, r_p)$. The algorithm must come across a ball B such that $5\rho\delta + r_p \leq \mathrm{radius}(B) \leq 10\rho\delta + 2r_p$. Before reaching the final $coarse(p)$, we show that the algorithm should have stopped and reported B as the final coarse ball, which would lead to contradiction.

Let x and y be points on S_δ^+ and S_δ^-, respectively, such that $\tilde{x} = \tilde{y} = \tilde{p}$. As $r_p = O(\epsilon^\gamma)$ by Lemma 6.1, $r_p \to 0$ as $n \to \infty$. As $\delta < 1/(1600\rho^2)$ and for sufficiently large n,

$$\mathrm{radius}(B) \leq 10\rho\delta + 2r_p < (f(\tilde{p}) - \delta)/2.$$

Thus for any points $a, b \in B \cap S_\delta^+$, $d(a,b) \leq 2\,\mathrm{radius}(B) < f(\tilde{p}) - \delta$. Let d_x be the maximum distance between $B \cap S_\delta^+$ and T_x. As $r_p < 1/(320\rho)$ for large n by Lemma 6.1 and $40\rho \leq 1/\sqrt{\delta}$, Lemma 2.2(i) implies that

$$\begin{aligned}
d_x &\leq (2\,\mathrm{radius}(B))^2 \\
&\leq (20\rho\delta + 4\,r_p)^2 \\
&\leq (\sqrt{\delta}/2 + 4\,r_p)^2 \\
&\leq \delta/4 + (4\sqrt{\delta} + 16\,r_p)r_p \\
&\leq \delta/4 + (\frac{8}{40\rho} + \frac{1}{20\rho})r_p \\
&\leq \delta/4 + \frac{r_p}{4\rho}.
\end{aligned}$$

Similarly, the maximum distance d_y between $B \cap S_\delta^-$ and T_y is at most $\delta + \frac{r_p}{4\rho}$. Therefore, if we enclose the noisy samples in $B \cap \mathbf{S}$ using a slab L parallel to T_p, then $\mathrm{width}(L) \leq d_x + d_y + 2\delta \leq 4\delta + \frac{r_p}{2\rho} \leq \mathrm{radius}(B)/(2\rho)$. Since we compute a 2-approximation of the thinnest slab, the slab width computed by the algorithm for B is at most $\mathrm{radius}(()B)/\rho$. But then, the algorithm should have stopped and reported B. This is a contradiction.

Next, we lower bound R_p. Clearly, $R_p \geq r_p$ by construction. We show that $R_p \geq 2\sqrt{\rho}\delta$. Let $B = B(p, R_p/\sqrt{\rho})$ and let $X = B(\tilde{p}, \delta)$. Note that $p \in X$ and X is tangent to S_δ^+ and S_δ^-. Lemma 2.1 implies that X is sandwiched between S_δ^+ and S_δ^- as $f(\tilde{p}) - \delta > \delta$.

Assume to the contrary that $R_p < 2\sqrt{\rho}\delta$. Then $\mathrm{radius}(B) < 2\delta$. Since $p \in X$, $B \cap X$ contains a ball with radius $\mathrm{radius}(B)/4$. The width of $\mathit{candidate}(p)$ is at most $R_p/\rho = \mathrm{radius}(B)/\sqrt{\rho}$. Thus, $(B \cap X) - \mathit{candidate}(s)$ contains a ball Y such that Y is empty, Y is sandwiched between S_δ^+ and S_δ^-, and

$$\begin{aligned}
\mathrm{radius}(Y) &\geq (\frac{1}{8} - \frac{1}{8\sqrt{\rho}}) \cdot \mathrm{radius}(B) \\
&\geq R_p/(16\sqrt{\rho}).
\end{aligned}$$

Since $R_p \geq r_p$ by construction, $\mathrm{radius}(Y) \geq r_p/(16\sqrt{\rho}) \geq r_p^2$ for sufficiently large n. By Lemma 6.1, $r_p^2 \geq \epsilon$. But then Y contains a noisy sample, a contradiction. Hence $R_p \geq \min\{2\sqrt{\rho}\delta, r_p\}$. $\square$

6.2 Rough normal estimate

We remark that the following analysis of the approximation error assumes that $r_p < 1$. Since $r_p = \mathrm{radius}(\mathit{initial}(p))^\gamma \leq O(\epsilon^\gamma)$ and $\epsilon = O((\frac{\ln^\omega n}{n})^{2/9})$, the smaller γ is, the larger n needs to be so that $r_p < 1$. This is a tradeoff between the running time and the number of samples needed to achieve a small approximation error. Of course, as γ is a constant with respect to n, γ does not affect the asymptotic behavior of the algorithm. Next, we show that $candidate(p)$ is a rough approximation of T_p in Lemma 6.4. To do this, we need one more technical lemma (Lemma 6.3), which proof is deferred to appendix.

Lemma 6.3 *Assume that $\delta \leq 1/2$. Let $p \in S_\alpha$ be a point. Let $d \leq 1/20$. Let H be a plane passing p and n_p. Then $S_\alpha \cap B(p,d) \cap H$ is a curve segment, and $S_\alpha \cap B(p,d)$ is a topological disk.*

Lemma 6.4 *Assume $\rho > 1$ and $\delta \leq 1/(1600\rho^2)$. For sufficiently large n, the angle between T_p and $candidate(p)$ is at most $20\rho\delta + 4r_p + 4/\rho$.*

Proof. Let $p \in S_\alpha$. If T_p and $candidate(p)$ are parallel, we are done. Suppose not. Let H be the plane spanned by n_p and the line perpendicular to $candidate(p)$. Figure 1 shows the cross-section on H. Let L denote the strip $candidate(p) \cap H$. Note that $\mathrm{width}(L) = \mathrm{width}(candidate(p))$. Let ℓ be

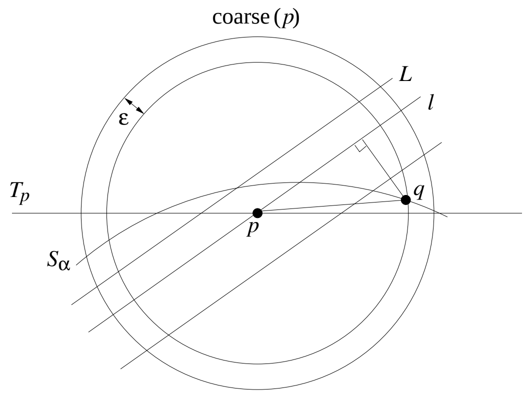

Figure 1: Illustration for the proof of Lemma 6.4.

the line on H through p parallel to L. By Lemma 6.3, the boundary of $B(p, R_p - \epsilon)$ intersects the curve $S_\alpha \cap H \cap coarse(p)$ at exactly two points. Let q be one of them.

Since $B(q,\epsilon)$ contains some noisy sample, $B(q,\epsilon)$ must intersect $candidate(p)$ as $candidate(p)$ contains all noisy samples inside $coarse(p)$. Thus $B(q,\epsilon) \cap H$ must intersect L.

Assume to the contrary that $\angle(T_p, \ell) > 20\rho\delta + 4r_p + 4/\rho$. Since $R_p < f(\tilde{p}) - \delta$, Lemma 2.2(i) implies that $\angle(pq, T_p) \leq \arcsin(R_p - \epsilon) \leq \arcsin R_p$, which is at most $\arcsin(10\rho\delta + 2r_p)$ by Lemma 6.2. Then $\angle(pq,\ell) \geq \angle(T_p,\ell) - \arcsin(10\rho\delta + 2r_p) > 20\rho\delta + 4r_p + 4/\rho - 20\rho\delta - 4r_p = 4/\rho$. It follows that

$$d(q,\ell) \geq d(p,q) \cdot \sin\angle(pq,\ell) > (R_p - \epsilon) \cdot \frac{2}{\rho} = \frac{R_p}{\rho} + \epsilon + \frac{R_p - 2\epsilon - \rho\epsilon}{\rho}.$$

As $r_p^2 \geq \epsilon$ by Lemma 6.1 and $\sqrt{\epsilon} \leq 1/(\rho+2)$ for large n by Lemma 4.1, we have $R_p \geq r_p \geq \sqrt{\epsilon} \geq (\rho+2)\epsilon$. This implies that $d(q,\ell) > R_p/\rho + \epsilon \geq \mathrm{width}(L) + \epsilon$, which contradicts the fact that $B(q,\epsilon) \cap H$ intersects L. $\square$

7 Refined neighborhood and point convergence

By Lemma 6.4, since $10\rho\delta \leq \pi/12$ as $\delta \leq 1/(1600\rho^2)$, the angle between T_p and $candidate(p)$ is at most $\pi/12$ for sufficiently large n. Since our algorithm enforces the angle between a tube axis ℓ and the direction orthogonal to $candidate(p)$ is at most $\pi/10$, $\angle(n_p, \ell) \leq \pi/12 + \pi/10 < \pi/5$. Recall that the radii of the tubes $L_p(\ell)$ and cylinders $C_p(\ell)$ are set to $r_p/4$. We want to show that REFINED NEIGHBORHOOD gives a provably good estimate of n_p. We need the following several technical lemmas, whose proofs are given in the appendix.

Lemma 7.1 *Let ℓ be an arbitrary line through a sample point p such that $\angle(n_p, \ell) \leq \pi/5$. For any $\alpha \in [-\delta, \delta]$, $L_p(\ell) \cap coarse(p) \cap S_\alpha$ is a topological disk inside $B(x, r_p/2)$, where $x = \ell \cap coarse(p) \cap S_\alpha$.*

The following lemma bounds some distances from p to S_α in terms of $\angle(n_p, \ell)$.

Lemma 7.2 *Let ℓ be an arbitrary line through a sample point p such that $\theta = \angle(n_p, \ell) \leq \pi/5$. For any α, let $p_1 \in S_\alpha$ with $\tilde{p_1} = \tilde{p}$. Let $x = \ell \cap coarse(p) \cap S_\alpha$. Then (i) $d(p_1, x) \leq \theta d(p, p_1)/2$, (ii) $d(p, x) \leq \frac{d(p,p_1)+\theta^2 d(p,p_1)^2/4}{\cos\theta}$, and (iii) $d(p, x) \geq \frac{d(p,p_1)-\theta^2 d(p,p_1)^2/4}{\cos\theta}$.*

The following lemma is for proving the point convergence later.

Lemma 7.3 *Let ℓ be a line through a sample point p such that $\theta = \angle(n_p, \ell) \leq \pi/5$. Let e be the segment $\ell \cap \mathbf{S} \cap coarse(p)$. Let q be the point $e \cap S$. Then $d(q, m) \leq 2\delta^2\theta^2/\cos\theta$, where m is the midpoint of e.*

Proof. Suppose $e = q_1 q_2$ where $q_1 \in S_\delta^+, q_2 \in S_\delta^-$. Let $p_1 p_2$ denote the normal segment passing through p, where $p_1 \in S_\delta^+, p_2 \in S_\delta^-$. Let $t_1 = \ell \cap T_{p_1}$ and $t_2 = \ell \cap T_{p_2}$. And let $t = \ell \cap T_{\tilde{p}}$. Refer to Figure 2. Note that t is the midpoint of the segment $t_1 t_2$. By Lemma 7.2 (ii) & (iii) and the fact that

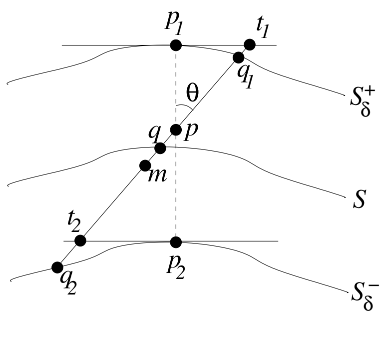

Figure 2: Illustration for the proof of Lemma 7.3.

$d(p, t_1) = d(p, p_1)/\cos\theta$, we get

$$|d(p, q_1) - d(p, t_1)| \leq \frac{\theta^2 d(p, p_1)^2}{4\cos\theta},$$

which is at most $\frac{\theta^2\delta^2}{\cos\theta}$ as $d(p, p_1) \leq 2\delta$. Since p, q_1, t are collinear, we conclude that

$$d(t_1, q_1) \leq |d(p, q_1) - d(p, t_1)| \leq \frac{\theta^2\delta^2}{\cos\theta}.$$

Similarly we can prove $d(t_2, q_2) \leq \delta^2\theta^2/\cos\theta$. Thus

$$d(t, m) \leq \delta^2\theta^2/\cos\theta. \tag{1}$$

The same argument also shows that $|d(p,q)-d(p,t)| \leq \frac{\theta^2\delta^2}{\cos\theta}$. Since p, q, t are collinear, we conclude that

$$d(t, q) \leq |d(p, q) - d(p, t)| \leq \frac{\theta^2\delta^2}{\cos\theta}. \tag{2}$$

Hence, by 1 and 2, $d(q, m) \leq d(t, m) + d(t, q) \leq 2\delta^2\theta^2/\cos\theta$. ◻

7.1 Height bounds of cylinders

We prove upper and lower bounds on the heights of cylinders tested by the algorithm.

Lemma 7.4 *Let $\theta = \angle(n_p, \ell) \leq \pi/5$. Assume that $\theta \leq \sqrt{r_p}$. Let p be the point $\ell \cap S \cap coarse(p)$. Let p^* be the center of $C_p(\ell)$. Then the height of $C_p(\ell)$ is at most $\frac{2\delta+\theta^2\delta^2}{\cos\theta} + 2r_p\sqrt{r_p}$; and $d(p^*, q) \leq \frac{2\theta^2\delta^2}{\cos\theta} + r_p\sqrt{r_p}$.*

Proof. Let $e = q_1q_2$ be the segment $\ell \cap \mathbf{S} \cap coarse(p)$, where $q_1 \in S^+_\delta, q_2 \in S^-_\delta$. Let B_1, B_2 be the upper and lower bases of $C_p(\ell)$ respectively. Assume that B_i lies on the same side of q as q_i. Let $b_1 = \ell \cap B_1$ and $b_2 = \ell \cap B_2$. Let p_1p_2 be the normal segment passing through p where $p_1 \in S^+_\delta, p_2 \in S^-_\delta$. Note that T_p, T_{p_1}, T_{p_2} are parallel. Refer to Figure 3

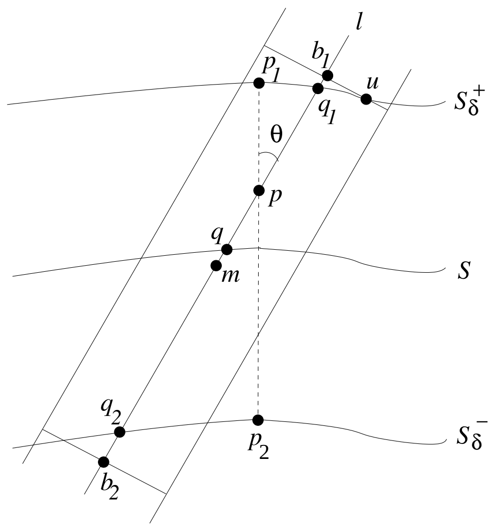

Figure 3: Illustration for the proof of Lemma 7.4.

We first show that $d(q_1, b_1) \leq r_p\sqrt{r_p}$. There are two cases to consider.

Case 1: b_1 lies between q_1 and q. Since $B(q_1, \epsilon)$ contains a sample point, we have $d(q_1, b_1) \leq \epsilon$, which is at most r_p^2 by Lemma 6.1.

Case 2: q_1 lies between b_1 and q. In this case, B_1 must intersect the patch $L_p(\ell) \cap S^+_\delta \cap coarse(p)$, otherwise we could have moved B_1 close to q to decrease the height of $C_p(\ell)$. Let u be a point at the intersection. So $d(q_1, b_1) \le r_p/4 \cdot \tan \angle q_1 u b_1$. We bound $\angle q_1 u b_1$ as follows.

$$\begin{aligned}\angle q_1 u b_1 &\le \angle(q_1 u, T_{q_1}) + \angle(n_{q_1}, \ell)\\ &\le \angle(q_1 u, T_{q_1}) + \angle(n_{p_1}, n_{q_1}) + \angle(n_{p_1}, \ell)\\ &= \angle(q_u, T_{q_1}) + \angle(n_{p_1}, n_{q_1}) + \theta.\end{aligned}$$

Since $d(q_1, u) \le r_p/2$ by Lemma 7.1, Lemma 2.2 (ii) implies that $\angle(q_1 u, T_{q_1}) \le \arcsin(r_p/2) \le r_p$. By Lemma 2.2 (ii), $\angle(n_{p_1}, n_{q_1}) \le 3d(p_1, q_1)$. By Lemma 7.2 (i), $d(p_1, q_1) \le \theta d(p, p_1)/2 \le \delta\theta$. So $\angle(n_{p_1}, n_{q_1}) \le 3\delta\theta$. Therefor $\angle q_1 u b_1 \le r_p + 3\delta\theta + \theta \le 2\sqrt{r_p}$ as $\theta \le \sqrt{r_p}$ by assumption. Hence $d(q_1, b_1) \le r_p/4 \cdot \tan(2\sqrt{r_p}) \le r_p\sqrt{r_p}$ for sufficiently small r_p.

Similarly we can prove $d(q_2, b_2) \le r_p\sqrt{r_p}$. It follows that

$$\begin{aligned}d(p, b_1) &\le d(p, q_1) + d(q_1, b_1)\\ &\overset{Lemma\ 7.2(ii)}{\le} \frac{d(p, p_1) + \theta^2 d(p, p_1)^2/4}{\cos\theta} + r_p\sqrt{r_p}.\end{aligned}$$

Similarly, we also have

$$d(p, b_2) \le \frac{d(p, p_2) + \theta^2 d(p, p_2)^2/4}{\cos\theta} + r_p\sqrt{r_p}.$$

Hence the height of $C_p(\ell)$ is $d(b_1, b_2) \le \frac{2\delta + \theta^2\delta^2}{\cos\theta} + 2r_p\sqrt{r_p}$.

Next, we upper bound $d(p^*, q)$. Let m be the midpoint of e. As p^* is the midpoint of $b_1 b_2$, we have $d(p^*, m) \le r_p\sqrt{r_p}$. On the other hand, by Lemma 7.3 (i), $d(q, m) \le 2\delta^2\theta^2/\cos\theta$. Thus by triangle inequality, we have

$$\begin{aligned}d(p^*, q) &\le d(q, m) + d(p^*, m)\\ &\le 2\delta^2\theta^2/\cos\theta + r_p\sqrt{r_p}.\end{aligned}$$

$\square$

Lemma 7.5 *Let $\theta = \angle(n_p, \ell) \le \pi/5$. Assume that $\theta > 32\sqrt{r_p}$. Then the height of $C_p(\ell)$ is at least* $\frac{2\delta - \theta^2\delta^2}{\cos\theta} + 2r_p\sqrt{r_p}$.

Proof. Assume that ℓ is vertical. Let H be the plane spanned by ℓ and n_p. Let p_1 be the point on S_δ such that $\tilde{p_1} = \tilde{p}$. Let ξ denote $S_\delta \cap L_p(\ell) \cap H$. By Lemma 6.3, ξ is a curve segment. Let a and b be the endpoints of ξ.

We first analyze the directions of the normals of ξ on H. Consider any point $z \in \xi$. Let n'_z be the projection of n_z onto H. Note that n'_z is the normal of ξ at z on H. Observe that $\angle(n_p, n'_z) \le \angle(n_p, n_z) = \angle(n_{p_1}, n_2)$. By Lemmas 7.1 and 7.2(i), $d(p_1, z) \le \theta\delta + r_p/2$. Lemma 2.2 (i) implies $\angle(n_{p_1}, n_z) \le \arcsin(\theta\delta + r_p/2) \le 2\theta\delta + r_p$. So $\angle(n_p, n'_z) \le 2\theta\delta + r_p$.

Assume that $\theta > 32\sqrt{r_p}$. We argue that ξ strictly increases or decreases from a to b. If n_p points to the left of ℓ, then n'_z also points to the left of ℓ, making an angle $(1 - 2\delta)\theta - r_p > 0$ as $\theta > 32\sqrt{r_p}$. So

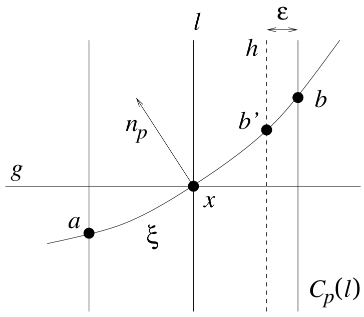

Figure 4: Illustration for the proof of Lemma 7.5.

ξ strictly increases from a to b. Conversely, if n_p points to the right of ℓ, then ξ strictly decreases from a to b.

W.l.o.g., assume that ξ strictly increases from a to b. Refer to Figure 4. Let q_1 be the point $\ell \cap S_\delta$. Let g be the horizontal line through q_1. Let h be the vertical line at distance ϵ from b. Let b' be the point $h \cap \xi$. Then

$$\begin{aligned} d(b', g) &\geq (r_p/4 - \epsilon)\tan((1-2\delta)\theta - r_p) \\ &\geq (r_p/4 - \epsilon)((1-2\delta)\theta - r_p). \end{aligned}$$

As $\epsilon \leq r_p^2$, $r_p/4 - \epsilon \geq r_p(1/4 - r_p)$ which is at least $r_p/8$ for sufficiently large n. Since $\theta > 32\sqrt{r_p}$, $(1-2\delta)\theta - r_p > 16\sqrt{r_p}$ when r_p is small. We conclude that $d(b', g) \geq 2r_p\sqrt{r_p}$.

Since the ball $B(b', \epsilon)$ contains a noisy sample, the upper end of $C_p(\ell)$ must be at distance at least $d(b', g) - \epsilon$ from g, which is at least $2r_p\sqrt{r_p} - \epsilon \geq r_p\sqrt{r_p}$ as $\epsilon \leq r_p^2$. By Lemma 7.2(iii), the distance from p to the upper end of $C_p(\ell)$ is at least $\frac{d(p,p_1) - \theta^2 d(p,p_1)^2/4}{\cos\theta} + r_p\sqrt{r_p}$. We can derive a similar lower bound for the distance from p to the lower end of $C_p(\ell)$. Hence the height of $C_p(\ell)$ is at least $\frac{2\delta - \theta^2\delta^2}{\cos\theta} + 2r_p\sqrt{r_p}$. ◻

7.2 Point convergence

We are now ready to show the refined cylinder direction ℓ_p^* and the cylinder center p^* approximate normal n_p and surface point $\tilde{p}$ well.

Lemma 7.6 (POINT CONVERGENCE) *Assume that $\delta \leq 1/(1600\rho^2)$. Let p^* and ℓ_p^* be the center point and approximate normal direction computed for p. Then $\angle(n_p, \ell_p^*) = O(\sqrt{r_p})$ and $d(\tilde{p}, p^*) = O(\sqrt{r_p})$.*

Proof. Consider the cylinder $C_p(\ell_1)$ computed by the algorithm. Let h be its height and let $\theta = \angle(n_p, \ell_p^*)$. Consider the caps covering the sphere of directions centered at p in REFINED NEIGHBORHOOD. n_p projects to a point inside some cap. Let ℓ' be the line through the center of this cap and p. So $C_p(\ell')$ is the corresponding cylinder. Let h' be its height and let $\theta' = \angle(n_p, \ell')$. Note that $\theta' \leq \sqrt{r_p}$.

Assume to the contrary that $\theta > 32\sqrt{r_p}$. By Lemmas 7.4 and 7.5, $h - h' \geq C - C'$, where

$$C = \frac{2\delta}{\cos\theta\cos\theta'}(\cos\theta' - \cos\theta), \text{ and}$$

$$C' = \frac{\delta^2}{\cos\theta\cos\theta'}(\theta^2\cos\theta' + {\theta'}^2\cos\theta) \le \frac{2\delta^2\theta^2}{\cos\theta\cos\theta'}.$$

Since $\theta \ge 32\theta'$,

$$\begin{aligned}
\cos\theta' - \cos\theta &= 2\sin\frac{\theta'+\theta}{2}\sin\frac{\theta-\theta'}{2} \\
&\ge 2\sin\frac{\theta}{2}\sin\frac{\theta}{4} \\
&\ge 2\cdot\frac{\theta}{4}\cdot\frac{\theta}{8} \\
&\ge \theta^2/16.
\end{aligned}$$

Thus,

$$\begin{aligned}
C - C' &\ge \frac{\delta\theta^2}{8\cos\theta\cos\theta'} - \frac{2\delta^2\theta^2}{\cos\theta\cos\theta'} \\
&\ge \frac{\delta\theta^2}{8\cos\theta\cos\theta'}(1-16\delta) \\
&> 0
\end{aligned}$$

as $\delta \le 1/(1600\rho^2)$. So $h - h' > 0$. This contradicts the fact that $C_p(\ell_p^*)$ has the minimum height. Hence $\theta \le 32\sqrt{r_p}$.

Let q be the point $\ell_p^* \cap S$. By Lemma 7.2 (i), $d(\tilde{p}, q) \le \theta d(p,\tilde{p})/2$, which is at most $\delta\theta$ as $d(p,\tilde{p}) \le 2\delta$. By Lemma 7.4, $d(p^*, q) \le \frac{2\theta^2\delta^2}{\cos\theta} + r_p\sqrt{r_p}$. Thus $d(\tilde{p}, p^*) \le d(\tilde{p}, q) + d(p^*, q) = r_p\sqrt{r_p} + O(\theta) = O(\sqrt{r_p})$ as $\theta \le 32\sqrt{r_p}$. ◻

8 Triangle normal and homeomorphism

Since the selected center points in C^* are not on the surface S, we cannot directly use the proofs in the cocone paper [3] for noiseless sample points. Instead we adapts their technique for our homeomorphism proof. We first need a concept called *Restricted Delaunay triangulation* on S, which was defined in say [13]. We will describe it first in the subsection below.

8.1 Restricted Delaunay triangulation

The *Voronoi diagram* V_{C^*} of C^* contains the *Voronoi cells* V_{p^*} for all points $p^* \in C^*$, where V_{p^*} is defined to be

$$\{x \in R^3 : d(x,p^*) \le d(x,q^*) \text{ for any } q^* \in C^* \text{ and } q^* \ne p^*\}.$$

The *Delaunay triangulation* D_{C^*} of C^* has a edge p^*q^* if $V_{p^*} \cap V_{q^*}$ is a face, has a triangle $p^*q^*r^*$ if $V_{p^*} \cap V_{q^*} \cap V_{r^*}$ is an edge, and has tetrahedron $p^*q^*r^*s^*$ if $V_{p^*} \cap V_{q^*} \cap V_{r^*} \cap V_{s^*}$ is a point.

The restriction of V_{C^*} to the surface S is called the *restricted Voronoi diagram* of C^* on S, denoted as $V_{C^*}|_S$. It contains all the *restricted Voronoi cells* $V_{p^*}|_S = V_{p^*} \cap S$. The dual of these restricted Voronoi cells defines the *restricted Delaunay triangulation* $D_{C^*}|_S$. Specifically, $D_{C^*}|_S$ has an edge

p^*q^* if and only if $V_{p^*}|_S \cap V_{q^*}|_S \neq \emptyset$, and a triangle $p^*q^*r^*$ if and only if $V_{p^*}|_S \cap V_{q^*}|_S \cap V_{r^*}|_S \neq \emptyset$. Assuming that S does not pass through a Voronoi vertex in general position, there is no tetrahedron in $D_{C^*}|_S$.

Edelsbrunner and Shah [13] showed that the underlying space of $D_{C^*}|_S$ is homeomorphic to S if the following *closed ball property* holds: each $V_{p^*}|_S$ is a topological disk, each nonempty pairwise intersection $V_{p^*}|_S \cap V_{q^*}|_S$ is a curve segment, and each nonempty triple intersection $V_{p^*}|_S \cap V_{q^*}|_S \cap V_{r^*}|_S$ is a single point.

8.2 Overview for homeomorphism

Now let's have an overview on how we proceed to prove the homeomorphism.

In the algorithm, all the center points are scanned to select a subset C^* of center points: when a center point p^* is selected, we delete all unselected center points u^* such that $d(p^*, u^*) \leq r_p^{1/8}$. We prove that the output manifold by running the cocone algorithm on C^* is a faithful reconstruction of S. We assume the points in C^* are in general position. In practice most Delaunay triangulation codes simulate general position, so this is not an unreasonable assumption.

In the noiseless case, Amenta, Choi, Dey, and Leekha [3] prove that the cocone algorithm returns a faithful reconstruction. The key step is proving that the normal of each triangle converges to the normal at some triangle vertex. Based on this result, one can then show that the restricted Delaunay triangles are not removed during the manifold extraction. This is important as it guarantees that there is some underlying manifold to be extracted. Furthermore, the convergence of triangle normals implies that the output surface is locally flat, which is instrumental to proving the homeomorphism.

The main tool in proving the homeomorphism and the convergence of normals of triangles is that there are two large tangent balls on opposite sides of S that touch S at each noiseless sample. The radii of these tangent balls are at least the local feature size of the sample, and these tangent balls do not contain other noiseless samples. We claim that an analogous result holds in the noisy case. So the proofs in [3] can be adapted here.

Specifically, let p^* be the center point computed for a noisy sample, and let $p^* \in S_\alpha$. We claim that there are two large tangent balls B_1 and B_2 on opposite sides of S_α that touch S_α at p^*. The radii of B_1 and B_2 are at least c for some constant $0 < c < 1$. Note that c is very large compared to the inter-center distances which is $O(r_p^{1/8})$ after pruning. We do not need the radii of B_1 and B_2 to be proportional to the local feature size because we use uniform sampling.

By the point convergence lemma, the selected center points C^* lie inside a thin shell. We will look at this and some basic properties of the selected centers points C^* in Section 8.3. Then in Section 8.4, we will prove that there are two large empty tangent balls at each center point p^*. With this lemma, we prove that the union of all restricted Delaunay triangles is homeomorphic to S in Section 8.6. In Section 8.5, we review the cocone reconstruction algorithm. Finally in Sections 8.7, 8.8 and 8.9, we prove the properties on the sizes and normals of the reconstructed triangles and the homeomorphism between the reconstruction and the original surface S.

8.3 Thin shell

The point convergence lemma (Lemma 7.6) immediately implies the following lemma, which says that the center points are close to the surface S, and the estimated normals are close to the surface normals.

Lemma 8.1 *For any center point p^*, $d(\tilde{p^*}, \tilde{p}) = O(\epsilon^{\gamma/2})$ and $\angle(n_{p^*}, n_p) = O(\epsilon^{\gamma/2})$. Moreover, $d(p^*, \tilde{p^*}) = O(\epsilon^{\gamma/2})$.*

Proof. By Lemma 7.6 and Lemma 6.1, $d(p^*, \tilde{p}) = O(\sqrt{r_p}) = O(\epsilon^{\gamma/2})$. As $\tilde{p^*}$ is the closest point on S to p^*, $d(p^*, \tilde{p^*}) \leq d(p^*, \tilde{p}) = O(\epsilon^{\gamma/2})$. Thus $d(\tilde{p^*}, \tilde{p}) \leq d(p^*, \tilde{p^*}) + d(p^*, \tilde{p}) = O(\epsilon^{\gamma/2})$. By Lemma 2.2 (ii), $\angle(n_{p^*}, n_p) \leq 3d(\tilde{p^*}, \tilde{p}) = O(\epsilon^{\gamma/2})$.

Moreover, $d(p^*, \tilde{p^*}) \leq d(p^*, \tilde{p}) = O(\epsilon^{\gamma/2})$. ◻

By this lemma, we have $d(p^*, \tilde{p^*}) = O(\epsilon^{\gamma/2})$ for any $p^* \in C^*$. Let $\Delta = \max_{\{p^* \in C^*\}} d(p^*, \tilde{p^*})$, which is $O(\epsilon^{\gamma/2})$. Let $\mathbf{S}_\Delta$ be the volume between S^+_Δ and S^-_Δ. Then all center points in C^* lie inside $\mathbf{S}_\Delta$. The following lemma says that the center points C^* in $\mathbf{S}_\Delta$ are dense.

Lemma 8.2 *For each point $q \in \mathbf{S}_\Delta$, there is a center point $p^* \in C^*$ such that $d(p^*, q) = O(\epsilon^{\gamma/8})$.*

Proof. The ball $B(q, \epsilon)$ contains a sample point $s \in P$ by our deterministic noise model. Then $d(q, s^*) \leq d(q, s) + d(s, \tilde{s}) + d(\tilde{s}, s^*) \leq \epsilon + (\Delta + \epsilon) + d(\tilde{s}, s^*) \leq O(\epsilon^{\gamma/2}) + d(\tilde{s}, s^*)$. Since $d(\tilde{s}, s^*) \leq d(\tilde{s}, \tilde{s^*}) + d(\tilde{s^*}, s^*) = O(\epsilon^{\gamma/2})$ by Lemma 8.1, we have $d(q, s^*) = O(\epsilon^{\gamma/2})$.

If $s^* \in C^*$, then we are done. Otherwise s^* is decimated when a center point p^* is selected. This means $d(s^*, p^*) \leq \sqrt[8]{r_p} = O(\epsilon^{\gamma/8})$. Thus $d(q, p^*) \leq d(q, s^*) + d(s^*, p^*) = O(\epsilon^{\gamma/8})$. ◻

Corollary 8.1 *For any $p^* \in C^*$, $V_{p^*}|_{S_alpha}$ is contained in a ball of radius $O(\epsilon^{\gamma/8})$ centered at p^*, where $\alpha \leq \Delta$.*

Next, we prove that after pruning, the distance between the selected center points is large when compared with Δ.

Lemma 8.3 *Let $p^*, q^* \in C^*$. Then $d(p^*, q^*) \geq \epsilon_2^{\gamma/8}$.*

Proof. Assume that, in the algorithm, p^* is selected before q^*. So $d(p^*, q^*) > r_p^{1/8}$. Recall that $r_p = \text{radius}(\textit{initial}(p))^\gamma$ and $\text{radius}(\textit{initial}(p)) \geq \epsilon_2$ by Lemma 6.1. So $d(p^*, q^*) \geq \epsilon_2^{\gamma/8}$. ◻

8.4 Twin empty tangent balls

We then prove that for any center point $p^* \in S_\alpha$, where $\alpha \leq \Delta$, there are twin balls tangent to S_α at p^* and empty of the center points in C^*. This lemma allows us to adapt many proofs from the cocone paper [3].

Lemma 8.4 *Suppose $p^* \in C^*$ is on S_α. For sufficiently large n, there are two balls tangent to S_α at p^* of radius $1/10$ on opposite sides of S_α such that their interiors do not contain any center point in C^*.*

Proof. We will consider only one side of S_α as the arguments for the other side are symmetric. Let m be a point at unit distance from $\tilde{p^*}$ in the normal direction. Let $M = B(m, 1-\Delta)$. Since $f_{\min} = 1$, the following holds: M is tangent to $\mathbf{S}_\Delta$, M is free from intersecting the interior of $\mathbf{S}_\Delta$, and M is empty of center points in C^*. Let q be the tangent point between M and S_Δ. We consider another ball $B = B(p^*, \epsilon_2^{\gamma/8})$, which is empty of center points in C^* by Lemma 8.3.

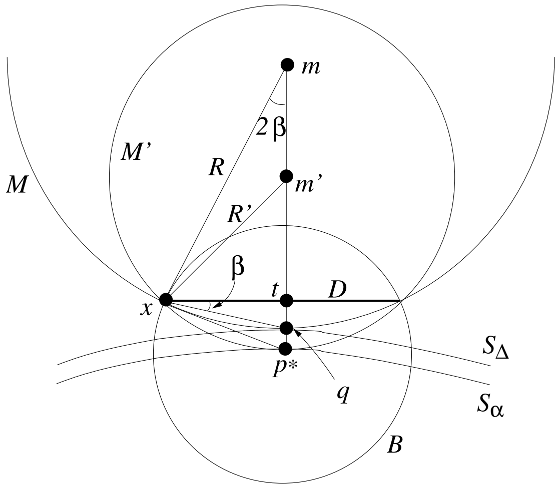

Figure 5: Illustration for the proof of Lemma 8.4.

Refer to Figure 5. Let D be the circle $\partial B \cap \partial M$. Let t be the center of D. Note that t lies on p^*m. We push M through circle D towards p^* until it touches p^* to produce a smaller ball $M' = B(m', r')$. We show that M' satisfies the lemma requirement by showing that $r' \geq 1/10$. Note that $m' \in p^*m$. So $M' \subset B \cup M$, which implies that M' does not contain any center point in C^*.

Consider a point x on the circle D. Note that $d(x, m') = r'$. As $\epsilon = O(\epsilon_2^{2/3})$, we have

$$d(x, q) \leq d(p^*, x) + d(p^*, q) \leq \epsilon_2^{\gamma/8} + 2\Delta.$$

Since $\Delta = O(\epsilon^{\gamma/2})$ and $\epsilon_2 = \Omega(\epsilon^{2/3})$, we have $2\Delta = O((\epsilon_2^{3/2})^{\gamma/2}) = O(\epsilon_2^{3\gamma/4})$, which is less than $\epsilon_2^{\gamma/8}$ for sufficiently large n. So $d(x, q) \leq 2\epsilon_2^{\gamma/8}$.

Let $\beta = \angle qxt$. Note that $\angle p^*mx = 2\beta$. Then $\sin\beta = \frac{d(x,q)/2}{1-\Delta} \leq \frac{\epsilon_2^{\gamma/8}}{1/2} = 2\epsilon_2^{\gamma/8}$. Thus

$$d(q, t) = d(x, q)\sin\beta \leq (2\epsilon_2^{\gamma/8})^2 = 4\epsilon_2^{\gamma/4}.$$

It follows that

$$\sin\angle p^*xt = \frac{d(p^*, q) + d(q, t)}{d(p^*, x)} \leq \frac{2\Delta + 4\epsilon_2^{\gamma/4}}{\epsilon_2^{\gamma/8}} = \frac{O(\epsilon_2^{3\gamma/4}) + 4\epsilon_2^{\gamma/4}}{\epsilon_2^{\gamma/8}}.$$

As $O(\epsilon_2^{3\gamma/4}) \leq \epsilon_2^{\gamma/4}$ for sufficiently large n,

$$\sin\angle p^*xt \leq \frac{5\epsilon_2^{\gamma/4}}{\epsilon_2^{\gamma/8}} = 5\epsilon_2^{\gamma/8}.$$

Hence $r' = \frac{d(p^*,x)/2}{\sin\angle p^*xt} \geq \frac{\epsilon_2^{\gamma/8}/2}{5\epsilon_2^{\gamma/8}} \geq 1/10.$ ◻

Consider a vertex $p^* \in S_\alpha$ of a restricted Delaunay triangle T. Since there are twin big balls tangent to S_α at p^* by the previous lemma, and the circumradius of a restricted Delaunay triangle is $O(\epsilon^{\gamma/8})$ by Corollary 8.1, we can prove the following lemma by using almost the same argument in Lemma 7 in [1]. The details of its proof is in the appendix.

Lemma 8.5 *Let τ be a restricted Delaunay triangle, and n_τ be the line normal to τ. Then $\angle(n_\tau, n_{p^*}) = O(\epsilon^{\gamma/8})$ for any vertex p^* of τ.*

With this lemma, by using almost the same arguments as Theorem 2 in [1], we have the following lemma. The details of its proof is in the appendix.

Lemma 8.6 *The union of restricted Delaunay triangles is homeomorphic to S.*

8.5 Review of the cocone algorithm

We will first review the cocone algorithm [3] because some understanding is needed to prove the homeomorphism.

First of all, the Delaunay triangulation of C^* is computed. For any point $p^* \in C^*$, let $e = w_1w_2$ be an edge in V_{p^*}. Suppose ℓ_p^* be the estimated normal by previous algorithm at p^*. Let $\theta = \pi/8$. We then compute $\angle(p^*w_1, \ell_p^*)$ and $\angle(p^*w_2, \ell_p^*)$, and check if the range of angles between these two angles intersect $[\pi/2-\theta, \pi/2+\theta]$. If it does, it actually means the Voronoi edge e intersects the complement of the double cone with apex p^* and making an angle of $\pi/2-\theta$ with axis ℓ_p^*, and we mark e. A Delaunay triangle T is put in W if its dual edge e is marked due to all three Voronoi cells defining e. Note that in the original algorithm [3], they estimate the surface normal with the *polar normal*, which is the direction from p^* to the farthest vertex in the Voronoi cell V_{p^*} at p^*. Since we have already obtained good normal estimates ℓ_p^* in the POINT ESTIMATION step, we don't need to estimate the polar normals.

After the candidate triangles are collected in W, the algorithm proceed to the manifold extraction step. First, all triangles incident to *sharp edges* are deleted. An edge is called *sharp* if the angle between any two consecutive triangles around the edge is more than $3\pi/2$, or the edge is incident to only one triangle. Next, the outer boundary N of the remained triangles is extracted by a depth-first walk along the outer boundary of each of its connected components.

It is proved in [3] that N is a manifold homeomorphic to the surface S. In order to establish that, they need three basic important conditions, which are:

(I) restricted Delaunay condition: W contains all restricted triangles;

(II) small triangle condition: the circumradius of each triangle in W is $O(\epsilon)$; and

(III) flat triangle condition: the normal to each triangle $T \in W$ makes a small angle $O(\epsilon)$ with the surface normal at the vertex p where p is the vertex with the largest interior angle in T.

As in our case, all the center points in C^* may not lie on S, the proofs of the original cocone paper [3] do not automatically apply. We have to generalize them so that they apply to samples points not on surface S as well. In the following sections, we will first prove the three basic conditions. Note that the small value $O(\epsilon)$ will be replaced by $O(\epsilon^{\gamma/16})$. And then we go forward to prove the homeomorphism.

8.6 Restricted Delaunay condition

We show that all restricted Delaunay triangles belong to W. So W contains triangles that form a manifold homeomorphic to S by Lemma 8.6. This is essential as otherwise there is no hope that the cocone algorithm will succeed in extracting a manifold from W.

Lemma 8.7 *Let v be any restricted Voronoi vertex in $V_{p^*}|_S$. Then $\angle(p^*v, n_p) \geq \pi/2 - O(\epsilon^{\gamma/8})$.*

Proof. Let $p^*q^*s^*$ be the restricted Delaunay triangle dual to v. Let B be the Delaunay ball centered at v and passing through p^*, q^* and s^*. As $2d(p^*, v) = 2\mathrm{radius}(B) \geq d(p^*, s^*) \geq \epsilon_2^{\gamma/8}$ by Lemma 8.3, we have $d(p^*, v) \geq d(p^*, s^*)/2 \geq \epsilon_2^{\gamma/8}/2$. On the other hand, as $v \in V_{p^*}|_S$, $d(p^*, v) = O(\epsilon^{\gamma/8})$ by Corollary 8.1. Also as $d(p^*, \tilde{p}) = O(\sqrt{r_p}) = O(\epsilon^{\gamma/2})$ by Lemma 7.6 and Lemma 6.1, we have $d(\tilde{p}, v) \leq d(p^*, \tilde{p}) + d(p^*, v) = O(\epsilon^{\gamma/8})$.

Now we have

$$\begin{aligned}
\angle p^*v\tilde{p} &\leq \arcsin \frac{d(p^*, \tilde{p})}{d(p^*, v)} \\
&= \frac{O(\epsilon^{\gamma/2})}{\epsilon_2^{\gamma/8}/2} \\
&= \frac{O(\epsilon^{\gamma/2})}{\Omega(\epsilon^{(3/2)(\gamma/8)})} \\
&= O(\epsilon^{5\gamma/16}).
\end{aligned}$$

By Lemma 2.2 (i), $\angle(\tilde{p}v, T_p) \leq \arcsin d(\tilde{p}, v) = O(\epsilon^{\gamma/8})$. So $\angle(\tilde{p}v, n_p) \geq \pi/2 - O(\epsilon^{\gamma/8})$. In all, $\angle(p^*v, n_p) \geq \angle(\tilde{p}v, n_p) - \angle p^*v\tilde{p} \geq \pi/2 - O(\epsilon^{\gamma/8}) - O(\epsilon^{5\gamma/16}) \geq \pi/2 - O(\epsilon^{\gamma/8})$. □

Lemma 8.8 *All restricted Delaunay triangles are in W.*

Proof. Let v be the restricted Voronoi vertex dual to the restricted Delaunay triangle $p^*q^*s^*$. By Lemma 8.7, $\angle(p^*v, n_p) \geq \pi/2 - O(\epsilon^{\gamma/8})$. Also $\angle(n_p, \ell_p^*) = O(\sqrt{r_p}) = O(\epsilon^{\gamma/2})$ by Lemma 7.6 and Lemma 6.1. The triangle inequality implies that $\angle(p^*v, \ell_p^*) \geq \angle(p^*v, n_p) - \angle(n_p, \ell_p^*) \geq \pi/2 - O(\epsilon^{\gamma/8}) - O(\epsilon^{\gamma/2}) \geq \pi/2 - O(\epsilon^{\gamma/8}) > \pi/2 - \pi/8$. Similarly, both $\angle(q^*v, n_q)$ and $\angle(s^*v, n_s)$ are greater than $\pi/2 - \pi/8$. Hence $p^*q^*s^*$ will be included in W by the algorithm. □

8.7 Small triangle condition

We prove that the circumcircle of each candidate triangle T in W is small; more precisely, its radius is at most $\epsilon^{\gamma/16}$.

Since there are twin big empty tangent balls tangent to S_α at center point $p^* \in S_\alpha$ by Lemma 8.4, we have the following lemma by using the same argument as in Lemma 5 in [1]. The details of this proof is included in the appendix.

Lemma 8.9 *Let v be any point in V_{p^*} such that $d(p^*, v) \geq \epsilon^{\gamma/16}$. Then $\angle(p^*v, n_{p^*}) = O(\epsilon^{\gamma/16})$.*

Lemma 8.10 *The radius of the smallest Delaunay ball of any triangle $T \in W$ is at most $\epsilon^{\gamma/16}$. Hence, the circumradius of T is at most $\epsilon^{\gamma/16}$.*

Proof. Let e be the dual edge of T, and p^* be a vertex of T. As the algorithm selects T, there is a point $v \in e$ so that $\angle(p^*v, \ell_p^*) \geq \pi/2 - \pi/8$. And by Lemma 7.6, $\angle(\ell_p^*, n_p) = O(\sqrt{r_p}) = O(\epsilon^{\gamma/2})$. Also by Lemma 8.1, $\angle(n_p, n_{p^*}) = O(\epsilon^{\gamma/2})$. Thus the triangle inequality implies that $\angle(p^*v, n_{p^*}) \geq \angle(p^*v, \ell_p^*) - \angle(\ell_p^*, n_p) - \angle(n_p, n_{p^*}) \geq \pi/2 - \pi/8 - O(\epsilon^{\gamma/2}) = \omega(\epsilon^{\gamma/16})$. Then by the contrapositive of Lemma 8.9, $d(p^*, v) < \epsilon^{\gamma/16}$. ◻

8.8 Flat triangle condition

We prove that the normal of each candidate triangle $T \in W$ makes a small angle of $O(\epsilon^{\gamma/16})$ radians with the surface normal at the vertex of T that has the largest vertex angle.

Since there are twin big empty balls tangent to S_α at center point $p^* \in S_\alpha$ by Lemma 8.4, and the candidate triangles in W are small by Lemma 8.10, we have the following theorem of triangle normal convergence by using the same argument in Theorem 11 in [3]. The details of this proof is in the appendix.

Lemma 8.11 *The normal to each candidate triangle $T \in W$ makes an angle of $O(\epsilon^{\gamma/16})$ radians with n_{p^*} where p^* is the vertex of T that has the largest vertex angle.*

This lemma says that the normal of triangle $T \in W$ is close to the surface normal at a vertex of T. As T is small by the small triangle condition, we can also show for any point x in the interior of T, $n_{\tilde{x}}$ is close to the normal of T.

Lemma 8.12 *Let x be any point on a triangle $T \in W$. Then*

(i) $d(x, \tilde{x}) = O(\epsilon^{\gamma/16})$, and

(ii) $\angle(n_{\tilde{x}}, n_T) = O(\epsilon^{\gamma/16})$, where n_T is the normal line to T.

Proof. Consider (i). Let p^* be a vertex of T. By Lemma 8.10, $d(x, p^*) \leq 2\epsilon^{\gamma/16}$, and by Lemma 7.6, $d(p^*, \tilde{p}) = O(\sqrt{r_p}) = O(\epsilon^{\gamma/2})$. So $d(x, \tilde{x}) \leq d(x, \tilde{p}) \leq d(x, p^*) + d(p^*, \tilde{p}) = O(\epsilon^{\gamma/16})$.

Consider (ii). By Lemma 8.1, $d(\tilde{p^*}, p^*) = O(\epsilon^{\gamma/2})$. By Lemma 8.10, $d(p^*, x) \leq 2\epsilon^{\gamma/16}$. By part (i), $d(x, \tilde{x}) = O(\epsilon^{\gamma/16})$. So $d(\tilde{p^*}, \tilde{x}) \leq d(\tilde{p^*}, p^*) + d(p^*, x) + d(x, \tilde{x}) = O(\epsilon^{\gamma/16})$. Thus by Lemma 2.2 (ii), $\angle(n_{p^*}, n_{\tilde{x}}) \leq 3d(\tilde{p^*}, \tilde{x}) = O(\epsilon^{\gamma/16})$. Also by Lemma 8.11, $\angle(n_T, n_{p^*}) = O(\epsilon^{\gamma/16})$. Hence $\angle(n_T, n_{\tilde{x}}) \leq \angle(n_T, n_{p^*}) + \angle(n_{p^*}, n_{\tilde{x}}) = O(\epsilon^{\gamma/16})$. ◻

8.9 Homeomorphism

A piecewise-linear manifold N is extracted from W in the manifold extraction step. It begins by recursively deleting any triangle in W adjacent to a sharp edge. Let W' be the remaining set of triangles. The lemma below proves that none of the restricted Delaunay triangles are deleted, and so W' contains a manifold homeomorphic to S.

Lemma 8.13 *No restricted Delaunay triangle has a sharp edge.*

Proof. First, since the restricted Delaunay triangles from a manifold homeomorphic to S, no edge of a restricted Delaunay triangle is incident to only one triangle. It remains to prove that no pair of restricted Delaunay triangles sharing an edge makes an acute angle.

Let T and T' be a pair of adjacent restricted Delaunay triangles sharing a common edge e. Let p^* be a vertex of e. Since T and T' are restricted Delaunay triangles, they have restricted Delaunay balls D, D' centered at restricted Voronoi vertices $v, v' \in S$ respectively.

Let H be the supporting plane of circle $\partial D \cap \partial D'$. Note that e lies on H and H separates T and T'. It is clear that $\angle(T, T') \geq \angle(T, H) + \angle(T', H)$ as H separates T and T'. So it suffices to lower bound $\angle(T, H)$ and $\angle(T', H)$.

Let n_T and $n_{T'}$ be the normal lines to T and T', respectively. By triangle inequality, $\angle(T, H) \geq \angle(n_p, T) - \angle(n_p, H)$. To lower bound $\angle(n_p, T)$, we need to upper bound $\angle(n_p, n_T)$. By Lemmas 8.1 and 8.5,

$$\angle(n_p, n_T) \leq \angle(n_{p^*}, n_p) + \angle(n_T, n_{p^*}) = O(\epsilon^{\gamma/2}) + O(\epsilon^{\gamma/8}) = O(\epsilon^{\gamma/8}).$$

So $\angle(n_p, T) \geq \pi/2 - O(\epsilon^{\gamma/8})$.

To upper bound $\angle(n_p, H)$, we need to lower bound $\angle(n_p, vv')$ as $vv' \perp H$. By Corollary 8.1 and Lemma 8.1, $d(v, \tilde{p}) \leq d(v, p^*) + d(p^*, \tilde{p}) = O(\epsilon^{\gamma/8}) + O(\epsilon^{\gamma/2}) = O(\epsilon^{\gamma/8})$. Then Lemma 2.2 (ii) implies that $\angle(n_v, n_p) \leq 3d(v, \tilde{p}) = O(\epsilon^{\gamma/8})$. So if we can lower bound $\angle(n_v, vv')$ then we are done. By Corollary 8.1, $d(v, v') \leq d(p^*, v) + d(p^*, v') = O(\epsilon^{\gamma/8})$. Then by Lemma 2.2 (i), $\angle(vv', T_v) \leq \arcsin d(v, v') = O(\epsilon^{\gamma/8})$. So $\angle(n_v, vv') = \pi/2 - \angle(vv', T_v) = \pi/2 - O(\epsilon^{\gamma/8})$. Hence

$$\angle(vv', n_p) \geq \angle(vv', n_v) - \angle(n_v, n_p) \geq \pi/2 - O(\epsilon^{\gamma/8}) - O(\epsilon^{\gamma/8}).$$

Thus $\angle(n_p, H) = \pi/2 - \angle(vv', n_p) = O(\epsilon^{\gamma/8})$, i.e., H is nearly parallel to n_p.

Then by triangle inequality, $\angle(T, H) \geq \angle(n_p, T) - \angle(n_p, H) \geq \pi/2 - O(\epsilon^{\gamma/8})$. Similarly we can prove that $\angle(T', H) \geq \pi/2 - O(\epsilon^{\gamma/8})$.

Now we have $\angle(T, T') \geq \angle(T, H) + \angle(T', H) \geq \pi - O(\epsilon^{\gamma/8})$, which is obtuse. So e is not sharp. □

Then in the algorithm, the outer boundary N of W' is extracted as the output manifold. In the lemma below, we prove that two adjacent triangles in N make an obtuse angle.

Lemma 8.14 *Every pair of adjacent triangles in N meets at their common vertex at an angle at least $\pi - O(\epsilon^{\gamma/16}) > \pi/2$.*

Proof. Suppose T and T' share a common vertex p^*. As $\angle(n_{p^*}, n_T) = O(\epsilon^{\gamma/16})$ and $\angle(n_{p^*}, n_{T'}) = O(\epsilon^{\gamma/16})$ by Lemma 8.11, we have

$$\angle(n_T, n_{T'}) \leq \angle(n_{p^*}, n_T) + \angle(n_{p^*}, n_{T'}) = O(\epsilon^{\gamma/16}).$$

So $\angle(T, T') \geq \pi - \angle(n_T, n_{T'}) \geq \pi - O(\epsilon^{\gamma/16}) > \pi/2$. □

Homeomorphism proof. Let $\mu : R^3 \to S$ be the map from each point in R^3 to the closest point on S. First we prove that the restriction of μ to N is well-defined.

Lemma 8.15 *$\mu : N \to S$ is well-defined.*

Proof. Suppose to the contrary that μ is not well-defined. Then there is a point x on a triangle in W such that x has more than one closest point to S. This means x is a point on the medial axis. So $d(x, S)$ is at least the local feature size of some point on S, and thus $d(x, S) \geq 1$. This contradicts Lemma 8.12(i) that $d(x, S) = O(\epsilon^{\gamma/16}) < 1$. ◻

Then we prove that $\mu : N \to S$ is a homeomorphism. A function is called a *homeomorphism* if it is continuous and bijective, and its inverse is also continuous. The approach is to first show that μ is well-behaved on center points C^*, and then extend the analysis to the interiors of the triangles in N.

Lemma 8.16 *The function $\mu : N \to S$ is continuous, and so is its inverse.*

Proof. By Lemma 8.10, every point $x \in N$ is within distance of $\epsilon^{\gamma/16}$ from a vertex $p^* \in C^*$, and thus is within a distance of $d(x, p^*) + d(p^*, \tilde{p}) = O(\epsilon^{\gamma/16})$ from $\tilde{p} \in S$ by Lemma 7.6. Note that the functions μ and μ^{-1} are continuous except at the medial axis of S. So the lemma follows from the fact that N and S are continuous and N and S avoid the medial axis. ◻

Lemma 8.17 *Assume $p^* \in C^*$ lies on S_α. Let m be the center of a medial ball M tangent to S_α at p^*. No candidate triangle in W intersects the interior of the segment p^*m.*

Proof. Let $M' = B(m', 1/10)$ be an empty ball tangent to S_α at p^* in Lemma 8.4. Note $m' \in p^*m$ as $d(p^*, m) = 1 - \Delta > 1/10$. Suppose to the contrary that a triangle T in W intersects the interior of p^*m at some point x. Then either $x \in m'm\backslash\{m\}$ or on $x \in$ int(p^*m'). Note that $\tilde{x} = \tilde{p^*}$. By Lemma 8.1 and Lemma 8.12(i), $d(p^*, x) \leq d(p^*, \tilde{p^*}) + d(\tilde{x}, x) = O(\epsilon^{\gamma/2}) + O(\epsilon^{\gamma/16}) = O(\epsilon^{\gamma/16})$.

If $x \in m'm\backslash\{m\}$, then $d(p^*, x) \geq 1/10 = \omega(\epsilon^{\gamma/16})$, which is a contradiction. Now suppose that $x \in$ int(p^*m'). In order to intersect int(p^*m'), T has to intersect M' as $p^*m' \subset M'$, and so does the smallest Delaunay ball D of T. Let H be the supporting plane of the circle $\partial M' \cap \partial D$. By adapting the proof for Lemma 16 of [3], we can derive the contradiction that H separates int(p^*m') and T.

On one side of H, M' is contained in D, and on the other, D is contained in M'. Let H^+ be the open halfspace, in which D contains M'. We first show that $p^* \notin H^+$. Since M' is empty of center points in C^*, T has to lie in H^+. Since D is Delaunay, p^* cannot lie in the interior of D. But since p^* lie on $\partial M'$, it therefore cannot lie in H^+. Next We show that $m' \notin H^+$ either. As $m' \in M'$, if $m' \in H^+$, m' has to lie in D. As $x \in T \subset D$, $xm' \subset D$. Since $2 \cdot \mathrm{radius}(D) \geq d(x, m') \geq d(p^*, m') - d(p^*, x) \geq 1/10 - O(\epsilon^{\gamma/16})$, we have $\mathrm{radius}(D) \geq 1/20 - O(\epsilon^{\gamma/16}) > \epsilon^{\gamma/16}$, which contradicts Lemma 8.10. In all, $p^*, m' \notin H^+$, which implies H separates int(p^*m') and T. ◻

Corollary 8.2
(i) Every center point in C^* is a vertex of N.
(ii) The function μ is one-to-one from N to $\mu(p^*)$ for every $p^* \in C^*$.

Our homeomorphism proof proceeds in three short steps. We show that μ induces a homeomorphism on each triangle, then on each pair of adjacent triangles, and finally on N as a whole.

Now with Lemma 8.12(ii), by using the same arguments as in Lemma 18 in [3], we thus have the following lemma, which shows that μ induces a homeomorphism on each triangle, then on union of adjacent triangles around a vertex.

Lemma 8.18 *Let U be a region contained within one triangle $T \in N$ or in adjacent triangles of N. The function μ defines a homeomorphism between U and $\mu(U) \subset S$.*

Then by Corollary 8.2 and Lemma 8.18 and using same arguments as Theorem 19 in [3], we have that μ is a homeomorphism for the whole N.

Lemma 8.19 *The mapping μ defines a homeomorphism from the triangulation N to the surface S.*

Hence by the convergence lemmas (Lemmas 7.6, 8.11) , the homeomorphism lemma above (Lemma 8.19) and the lemma about the running time analysis (Lemma 4.2), we finally can get our main theorem for this chapter (Theorem 5.1).

9 Summary

The result in this paper successfully extends our result of curve reconstruction from noisy samples [7] to surface reconstruction from noisy samples with theoretical guarantees. Instead of using probabilistic sampling model as in curve case, we use deterministic sampling model, which makes the proofs simpler and more intuitive. The algorithms runs in $O(n^{2+\gamma})$ time where $0 < \gamma \leq 1/8$ is a constant. The output piecewise linear manifold is homeomorphic to S, is close to S in terms of Hausdorff distance, and has normals close to the surface normals on S. This result, together with the recent result by Dey and Goswami [10] are the first two results to do surface reconstruction from noisy samples with theoretical guarantees.

Acknowledgment

We thank the anonymous referees for suggestions that greatly improve the presentation of this paper. We also thank Sunil Arya, Beifang Chen, Mordecai Golin, Tamal Dey, and Edgar Ramos for valuable comments.

References

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

[2] N. Amenta, M Bern and M. Kamvysselis. A new Voronoi-based surface reconstruction algorithm. *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*, 1998, 415–421.

[3] N. Amenta, S. Choi, T. Dey and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. *International Journal on Computational Geometry & Applications*, 12, 2002, 125–141.

[4] N. Amenta, S. Choi and R.K. Kolluri. The power crust. *Proceedings of the 6th ACM Symposium on Solid Modeling and Applications*, 2001, 249–260.

[5] N. Amenta, S. Choi and R.K. Kolluri. The power crust, unions of balls, and the medial axis transform. *Computational Geometry: Theory and Applications*, 19, 2001, 127–153.

[6] J.-D. Boissonnat and F. Cazals. Smooth surface reconstruction via natural neighbour interpolation of distance functions. *Computational Geometry*, 22, 2002, 185–203.

[7] S.-W. Cheng, S. Funke, M. Golin, P. Kumar, S.-H. Poon, and E.A. Ramos. Curve reconstruction from noisy samples. *Proceedings of the 19th Annual Symposium on Computational Geometry*, 2003, 302–311.

[8] B. Curless and M. Levoy. A volumetric method for building complex models from range images. *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*, 1996, 303–312.

[9] T.K. Dey, J. Giesen, S. Goswami, J. Hudson, R. Wenger, and W. Zhao. Undersampling and oversampling in sample based shape modeling. *Proceedings of IEEE Visualization*, 2001, 83–90.

[10] T.K. Dey and S. Goswami. Provable surface reconstruction from noisy samples. *Proceedings of the 20th Annual Symposium on Computational Geometry*, 2004, 330–339.

[11] H. Edelsbrunner. Surface reconstruction by wrapping finite point sets in space. *Ricky Pollack and Eli Goodman Festschrift*, ed. B. Aronov, S. Basu, J. Pach and M. Sharir, Springer-Verlag, 2003, 379-404.

[12] H. Edelsbrunner and E.P. Mucke. Three-dimensional alpha shapes. *ACM Transactions on Graphics*, 13, 1994, 43-72.

[13] H. Edelsbrunner and N. Shah. Triangulating topological spaces. *International Journal on Computational Geometry and Applications*, 7, 1997, 365–378.

[14] S. Funke and E. A. Ramos. Smooth-surface reconstruction in near-linear time. *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithm*, 2002, 781–790.

[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] M. Habib, C. McDiarmid, J. Ramirez-Alfonsin and B. Reed. *Probabilistic Methods for Algorithmic Discrete Mathematics*, Springer Verlag Belin Heidelberg, 1998, 198-200.

[17] S. Har-Peled and K.R. Varadarajan. Approximate shape fitting via linearization. *Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science*, 2001, 66–73.

[18] H. Hoppe. Surface reconstruction from unorganized points. *Ph.D. Thesis*, Computer Science and Engineering, University of Washington, 1994.

[19] H. Hoppe, T. DeRose, T. Duchamp, H. Jin, J. McDonald, and W. Stuetzle. Piecewise smooth surface reconstruction. *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques*, 1994, 19–26.

[20] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Surface reconstruction from unorganized points. *Proceedings of the 19th Annual Conference on Computer Graphics and Interactive Techniques*, 1992, 71–78.

[21] G. Medioni and C.K. Tang. Inference of integrated surface, curve, and junction descriptions from sparse 3-D data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20, 1998, 1206–1223.

[22] W.S. Massey. *Algebraic Topology: An Introduction*, Springer-Verlag, Graduate texts in Mathematics 56, 1967.

[23] N.J. Mitra and A. Nguyen. Estimating surface normals in noisy point cloud data. *Proceedings of the 14th Annual ACM Symposium on Computational Geometry*, 2003, 322–328.

[24] C.K. Tang and G. Medioni. A curvature-augmented tensorial framework for shape inference from noisy, 3-D data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24, 2002, 858–864.

Appendix

Proof of Lemma 2.1

Let M_α be the medial ball of S_α touching a point $p \in S_\alpha$. By the definition of S_α, there is a medial ball M of S touching $\tilde{p}$ such that M and M_α have the same center and $\mathrm{radius}(M_\alpha) = \mathrm{radius}(M) - \alpha \geq f(\tilde{p}) - \alpha$. Let B be a ball of radius $f(\tilde{p}) - \alpha$ that touches S_α at p. If S_α intersects the interior of B, the medial axis of S_α intersects the interior of B. But then $\mathrm{radius}(M_\alpha) < \mathrm{radius}(B) = f(\tilde{p}) - \alpha$, a contradiction. $\square$

The next lemma consists of a lemma by Giesen and Wagner [15] and a lemma by Amenta and Bern [1].

Lemma 9.1 *Let $p, q \in S$ be two points.*

(i) If $d(p,q) = \lambda f(p)$ for some $\lambda < 1$, then $d(q, T_p) \leq \frac{\lambda^2}{2} f(p)$ and $\angle(pq, T_p) \leq \arcsin(\frac{\lambda}{2})$.

(ii) If $d(p,q) = \lambda \min\{f(p), f(q)\}$ for some $\lambda \leq 1/3$, then $\angle(n_p, n_q) \leq \frac{\lambda}{1-3\lambda}$.

We are ready to prove Lemmas 2.2 and 2.3.

Proof of Lemma 2.2

Let f_α denote the local feature size function for any point on S_α. As S and S_α share the same medial axis, $f_\alpha(p) \geq f(\tilde{p}) - \delta \geq 1/2$ for any point $p \in S_\alpha$.

Consider (i). We have $d(p,q) = \beta f_\alpha(p)$ for some $\beta \leq \lambda$. Lemma 9.1(i) implies that

$$\begin{aligned} d(q, T_p) &\leq \beta^2 f_\alpha(p)/2 \\ &= d(p,q)^2/(2f_\alpha(p)) \\ &\leq d(p,q)^2. \end{aligned}$$

And $\angle(n_p, n_q) \leq \arcsin(\frac{\beta}{2}) \leq \arcsin(\frac{d(p,q)}{2f_\alpha(p)})$, which is at most $\arcsin(d(p,q))$.

Consider (ii). Let g denote $\min\{f(\tilde{p}) - \delta, f(\tilde{q}) - \delta\}$. We then have

$$d(p,q) = \beta \min\{f_\alpha(p), f_\alpha(q)\}$$

for some $\beta \leq \lambda \leq 1/5$. Lemma 9.1(ii) implies that

$$\begin{aligned} \angle(n_p, n_q) &\leq \beta/(1-3\beta) \\ &\leq d(p,q)/((1-3\beta)g). \end{aligned}$$

Note that $g \geq 1 - \delta \geq 5/6$ and $1 - 3\beta \geq 2/5$. So $(1-3\beta)g \geq 1/3$, which implies that $\angle(n_p, n_q) \leq 3d(p,q)$. $\square$

Proof of Lemma 2.3

Let r be the point $\tilde{p} + (q - \tilde{q})$. W.l.o.g., assume that $\angle \tilde{p}pr \leq \angle \tilde{p}rp$. By Lemma 2.2, $\angle p\tilde{p}r \leq 3d$. Therefore, $\angle \tilde{p}rp \geq \pi/2 - 3d/2$. By sine law, $d(p,r) = \frac{d(p,\tilde{p}) \cdot \sin \angle p\tilde{p}r}{\sin \angle \tilde{p}rp} \leq \frac{3\delta d}{\cos(3d/2)}$. Since $\delta \leq 1/10$ and $\cos(3d/2) \geq \cos(3/10) > 0.9$, we have $d(p,r) \leq d/2$. By triangle inequality,

$$\begin{aligned} d(p,q) &\leq d(q,r) + d(p,r) \\ &= d(\tilde{p},\tilde{q}) + d(p,r) \\ &\leq 3d/2. \end{aligned}$$

Similarly, $d(p,q) \geq d(q,r) - d(p,r) \geq d/2$. $\square$

We need the Lemmas 9.2, 9.3, and 9.4, to prove Lemma 3.3.

Lemma 9.2 *Let $p \in S_\alpha$ be a point. Let k and d be constants such that $k \geq 2$ and $kd \leq 1/20$. Let ℓ be a line such that $d(p,\ell) \leq d$ and $\angle(n_p, \ell) \leq \pi/4$. Then ℓ intersects $S_\alpha \cap B(p,kd)$ at exactly one point.*

Proof. Let $C = S_\alpha \cap B(p,kd)$. We first show that ℓ intersects C. Suppose not. Translate ℓ towards p until ℓ touches C at some point q. We claim that $q \in \mathrm{int}(C)$. By this claim, ℓ is tangent to S_α at q. Observe that n_q lies on the plane through q and orthogonal to ℓ. This implies that $\angle(n_p, n_q) \geq \pi/2 - \angle(n_p, \ell) \geq \pi/4$, which contradicts Lemma 2.2(ii). We now prove the claim. Suppose that $q \in \partial C$. Then q lies on the boundary of $B(p,kd)$ and between the medial balls at p. Refer to Figure 6. Let ℓ_p be the support line of n_p. Let ℓ' be the projection of ℓ onto the plane containing ℓ_p and q. Starting with Lemma 2.2(i), we have $\angle(pq, T_p) \leq \arcsin(kd) \leq 1/10$ as $kd \leq 1/20$. Also, $\angle(\ell_p, \ell') \leq \angle(\ell_p, \ell) \leq \pi/4$. Thus

$$\begin{aligned} \angle(pq, \ell') &= \pi/2 - \angle(\ell_p, \ell') - \angle(pq, T_p) \\ &\geq \pi/4 - 1/10. \end{aligned}$$

But then $d(p, \ell) \geq d(p, \ell') = kd \sin \angle(pq, \ell') > d$ as $k \geq 2$ and $\sin(\pi/4 - 1/10) > 1/2$, a contradiction.

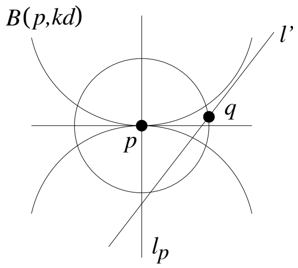

Figure 6: Illustration for the proof of Lemma 9.2.

Next, we show that ℓ intersects C at exactly one point. Suppose that there are two points $x, y \in \ell \cap C$. Note that $d(x, y) \leq 2kd$. By Lemma 2.2(i), $\angle(xy, T_x) \leq \arcsin(2kd) < 1/5$ as $kd \leq 1/20$. So $\angle(n_x, \ell) \geq \pi/2 - 1/5$, which implies that $\angle(n_x, n_p) \geq \pi/4 - 1/5 > 3kd$ as $kd \leq 1/20$. This contradicts Lemma 2.2(ii). $\square$

Lemma 9.3 *Assume that $\delta \leq 1/2$. Let $p \in S_\alpha$ be a point. Let $d \leq 1/20$. Let Y be the infinite cylinder with radius d and axis aligned with n_p. Let C be the connected component of $Y \cap S_\alpha$ containing p. Let H be a plane passing p and n_p. Then $C \cap H$ is a curve segment, and C is a topological disk.*

Proof. First, we show that $C \subset S_\alpha \cap B(p, 2d)$. Consider any point q on C. Observe that the two balls of radius $1 - \delta$ tangent to S_α at p do not intersect $S_\alpha \backslash \{p\}$. As d is less than the radius $1 - \delta$ of the two balls, it is clear that $\angle(pq, T_p) < \pi/4$, which implies $d(p, q) < 2d$.

By Lemma 9.2, ℓ intersects $S_\alpha \cap B(p, 2d)$ at exactly one point. Since $C \subset S_\alpha \cap B(p, 2d)$, ℓ intersects C with at most one point. Suppose ℓ does not intersect C. Then we can translate ℓ toward n_p. At some point, ℓ will be tangent to some point z on C, which means $n_z \perp n_p$. However, we can prove that n_z and n_p make an acute angle, which then leads to a contradiction. As $d(p, z) \leq 2d$ and by Lemma 2.2(ii), $\angle(n_p, n_z) \leq 3d(p, z) \leq 6d < \pi/2$. So $\ell \cap C$ is exactly one point. Therefore $C \cap H$ is a curve segment, and C is a topological disk. $\square$

Lemma 9.4 *Let p be a point on S_α. Let H be a plane passing through p making an acute angle with n_p. Let curve $F_\alpha = S_\alpha \cap H$. Then the radius of curvature of F_α at p is at least $(f(\tilde{p}) - \delta) \cos \angle(n_p, H)$.*

Proof. By Lemma 2.1, there are two tangent balls B_1, B_2 of S_α at p of radius $f(\tilde{p}) - \delta$ that do not intersect $S_\alpha \backslash \{p\}$. We show that there are two tangent disks D_1, D_2 of F_α at p of radius $(f(\tilde{p}) - \delta) \cos \angle(n_p, H)$ that do not intersect $F_\alpha \backslash \{p\}$.

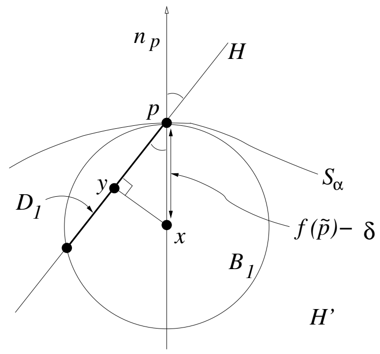

Figure 7: Illustration for the proof of Lemma 9.4.

Consider the plane H' containing n_p and the normal to H. Consider the cross-section of everything on H' in Figure 7. Let $D_1 = B_1 \cap H$. As B_1 does not intersect $S_\alpha \backslash \{p\}$, D_1 does not intersect $F_\alpha \backslash \{p\}$. Let x and y be the centers of B_1 and D_1, respectively. As $xy \perp py$, $\text{radius}(D_1) \geq d(p, y) = (f(\tilde{p}) - \delta) \cos \angle xpy = (f(\tilde{p}) - \delta) \cos \angle(n_p, H)$. ◻

Now we can do the proof for Lemma 3.3.

Proof of Lemma 3.3

First we establish the lower bound. Let H be a plane that contains the normal to S_α at p. Consider any point $q \in S_\alpha \cap Y \cap H$. Let C be the connected component of $S_\alpha \cap Y$ containing p. Let D be the projection area of C onto T_p. As C is a topological disk by Lemma 9.3, D is a disk. Thus $\text{area}(C) \geq \text{area}(D) \geq \pi r^2$.

Then we will establish the upper bound. The radius of curvature at any point q on C is at least $1 - \delta$. Consider the projection of C onto T_p. Let J be the curve segment $H \cap C$ according to Lemma 9.3. Refer

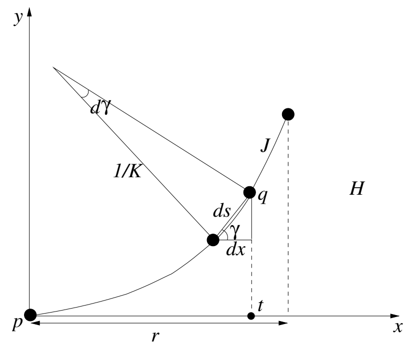

Figure 8: Illustration for the proof of Lemma 3.3.

to Figure 8: $H \cap T_p$ is the x-axis, and n_p is the y-axis. Let q be any point on J, and let t be its projection onto T_p. By Lemma 2.1, there are two tangent disks of J at p on H of radius $R = 1 - \alpha$ at p. As $r \leq R$, it is clear that $\angle(pq, T_p) \leq \pi/4$, which implies $d(q, t) \leq d(p, t) \leq r$. Thus $d(p, q) \leq \sqrt{2}r < 3r/2$. Then by Lemma 2.3, $d(\tilde{p}, \tilde{q}) \leq 2d(p, q) \leq 3r$. By Lemma 2.2(ii), $\angle(n_p, n_q) \leq 9r \leq 9/20 < \pi/6$, which implies $\angle(n_q, H) < \pi/6$ as n_p lies on H. By Lemma 9.4, the radius of curvature of J at q is at least $(1-\delta)\cos\angle(n_q, H) > (1-\delta)\cos\pi/6 = \frac{\sqrt{3}(1-\delta)}{2}$, which means the curvature $k(q)$ of any point on J is at most $\frac{2}{\sqrt{3}(1-\delta)}$. Let γ be the angle between the tangent of J at q and the horizontal. Let s be the arc length along J from p to q. Then as $ds = (\frac{1}{k(q)})d\gamma$,

$$d\gamma = k(q)ds = k(q)\sec\gamma dx \le \frac{2}{\sqrt{3}(1-\delta)}\sec\gamma dx.$$

So $\frac{\sqrt{3}(1-\delta)}{2}\cos\gamma d\gamma = dx$. Thus by doing integration on both sides, $\frac{\sqrt{3}(1-\delta)}{2}\sin\gamma \le x$, which implies $\sin\gamma \le \frac{2x}{\sqrt{3}(1-\delta)} \le \frac{2r}{\sqrt{3}(1-\delta)}$ and $\cos\gamma \ge \sqrt{\frac{2}{3}}$ as $r \le \frac{1-\delta}{20}$. Therefore, $ds = \sec\gamma dx \le \sqrt{\frac{3}{2}}dx$.

Suppose H' is the plane obtained by rotating H through an angle $d\theta$ around n_p (i.e. the y-axis in Figure 8 and Figure 9). Let J' be the curve segment $H' \cap C$. Suppose Y' be a smaller cylinder concentric to Y and passing through a point q on J. Let q' be the intersection point $J' \cap Y'$. Refer to Figure 9. The arc length dt from q to q' along the boundary of $Y' \cap C$ is approximately $d(q, q')$.

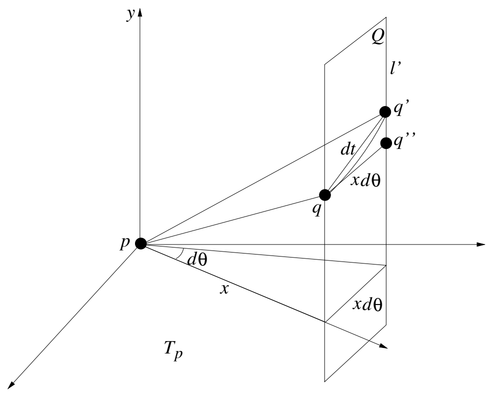

Figure 9: Illustration for the proof of Lemma 3.3.

Consider another plane Q passing through q and q' and perpendicular to T_p. Let ℓ' be the line through q' perpendicular to T_p. Let q'' be the projection of q onto ℓ'. Now $dt \approx d(q, q') = d(q, q'')\sec\angle q'qq'' = xd\theta\sec\angle q'qq''$. By Lemma 2.2 (i) and since dt is very small,

$$\begin{aligned}
\angle q'qq'' &= \angle(qq', T_p) \\
&\le \angle(qq', T_q) + \angle(T_q, T_p) \\
&\le \arcsin(d(q, q')) + \angle(n_p, n_q) \\
&\le 2dt + \angle(n_p, n_q) \\
&< \pi/6.
\end{aligned}$$

Thus $dt \le xd\theta\sec\frac{\pi}{6} = \frac{2}{\sqrt{3}}xd\theta$.

Refer to Figure 10. The surface area of C is $\int\int dsdt \le \int\int(\sqrt{\frac{3}{2}}dx)(\frac{2}{\sqrt{3}}xd\theta) \le \int_0^r\int_0^{2\pi}\sqrt{2}xd\theta dx = \int_0^r 2\sqrt{2}\pi xdx = [\sqrt{2}\pi x^2]_0^r = \sqrt{2}\pi r^2$. $\square$

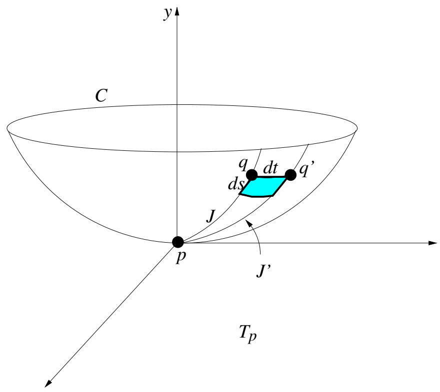

Figure 10: Illustration for the proof of Lemma 3.3.

Proof of Lemma 4.1

First, we claim that $\epsilon_2 = O(\sqrt[3]{\frac{\ln^\omega n}{n}})$. By this claim, since $\epsilon_2 = \Omega(\epsilon^{3/2})$ by the model, we have $\epsilon = O(\epsilon_2^{2/3}) = O((\frac{\ln^\omega n}{n})^{2/9})$. Then the upper bound on ϵ_1 follows as $\epsilon_1 = O(\epsilon)$ by the model.

Take a maximal set K of disjoint balls with radii $\epsilon_2/2$ and centers inside **S**. So all points in K are within a distance of $\delta + \epsilon_2/2$ from S. A packing argument shows that $|K| \leq \frac{(2\delta+\epsilon_2)\mathrm{area}(S)}{(4/3)\pi(\epsilon_2/2)^3} = O((\delta + \epsilon_2)/\epsilon_2^3)$. If we double the sizes of the balls in K, the expanded balls cover **S** entirely. Each expanded ball has radius ϵ_2 and contains at most $\ln^\omega n$ samples by the model. It follows that $|K| \geq n/\ln^\omega n$. Combining the upper and lower bounds for $|K|$, we get $n/\ln^\omega n = O((\delta+\epsilon_2)/\epsilon_2^3)$. If $\delta \leq \epsilon_2$, replacing δ by ϵ_2 yields $\epsilon_2 = O(\sqrt{\ln^\omega n/n})$. If $\delta > \epsilon_2$, as $\delta < 1$, replacing $\delta+\epsilon_2$ by 2 yields $\epsilon_2 = O(\sqrt[3]{\ln^\omega n/n})$. This proves our first claim.

Second, we claim that $\epsilon = \Omega(\sqrt{1/n})$. By this claim, since $\epsilon_2 = \Omega(\epsilon^{3/2})$ by the model, we have $\epsilon_2 = \Omega(1/n^{3/4})$.

Take a maximal set L of disjoint balls with radii ϵ and centered at points on S. Since each ball contains at least one sample by the model, $n \geq |L|$. By Lemma 3.3, the area of the intersection between S and each ball in L is $O(\epsilon^2)$. It follows that $|L| = \Omega(\mathrm{area}(S)/\epsilon^2) = \Omega(1/\epsilon^2)$. Combining the upper and lower bounds for L, we get $\epsilon = \Omega(1/\sqrt{n})$. This proves our second claim. ◻

Proof of Lemma 6.3

Let ℓ be the support line of n_p. It suffices to show that as we translate ℓ on H to a distance d from p, $\ell \cap S_\alpha$ traces a single curve segment inside $B(p,d)$. By Lemma 9.2, $\ell \cap S_\alpha$ traces a single curve segment ξ inside $B(p,2d)$. Assume to the contrary that ξ exits $B(p,d)$ at a point x, and reenters $B(p,d)$ at a point y. Let $\xi(x,y)$ denote the subcurve between x and y that lies outside $B(p,d)$. By Lemma 2.2(i), $\angle(px, T_p) \leq \arcsin(d) \leq 2d$. Similarly, $\angle(py, T_p) \leq 2d$. Thus the vector n from p to the midpoint of xy makes an angle at least $\pi/2 - 2d$ with n_p. Observe that n is parallel to a curve normal of some point $z \in \xi(x,y)$. It follows that $\angle(n_z, n_p) \geq \angle(n, n_p) \geq \pi/2 - 2d$. But this contradicts Lemma 2.2(ii). ◻

Proof of Lemma 7.1

Let $C = L_p(\ell) \cap S_\sigma \cap coarse(p)$. It follows from Lemma 9.2 that C is a topological disk, and for any line $h \in L_p(\ell)$, since $d(p,h) \le r_p/4$, h intersects $S_\sigma \cap B(x, r_p/2)$ exactly once. Hence $C \subseteq B(x, r_p/2)$. $\square$

We then prove Lemma 7.2, which bounds some distances from p to S_α in terms of $\angle(n_p, \ell)$.

Proof of Lemma 7.2

Refer to Figure 11. Let h be the line through p_1 parallel to ℓ. Let Y be the infinite tube with axis h and p on its boundary. So the radius of Y is $d(p,p_1)\sin\theta$. Then as in the proof of Lemma 7.1, $d(p_1,x) \le \mathrm{radius}(Y)/2$, which is at most $d(p,p_1)/2 \cdot \sin\theta \le \theta d(p,p_1)/2$.

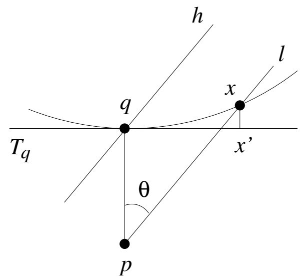

Figure 11: Illustration for the proof of Lemma 7.2.

Let x' be the orthogonal projection of x onto T_{p_1}. By Lemma 2.2(i), $d(x,x') \le d(p_1,x)^2$, which is at most $\theta^2 d(p,p_1)^2/4$. Thus

$$\begin{aligned} d(p,x) &\le \frac{d(p,p_1) + d(x,x')}{\cos\theta} \\ &\le \frac{d(p,p_1) + \theta^2 d(p,p_1)^2/4}{\cos\theta}. \end{aligned}$$

Symmetrically, $d(p,x) \ge \frac{d(p,p_1) - d(x,x')}{\cos\theta}$ gives the lower bound. $\square$

Proof of Lemma 8.5

Let s^* be the vertex of τ that has the largest vertex angle ($\ge \pi/3$). Suppose that s^* is on the surface S_α. Let C be the circumcircle of T. By Lemma 8.4, there are twin empty balls B, B' of radius $1/10$ tangent to S_α at s^*. These balls intersect the supporting plane of τ in twin disks D, D' tangent at s^* such that $D \subset B$ and $D' \subset B'$. Let R be the common radius of D and D'. See Figure 12. Our first upper bound R in terms of $\mathrm{radius}(C)$.

Since the twin tangent balls at s^* are empty, the twin disks do not contain the vertices of τ. In order to maximize R relative to $\mathrm{radius}(C)$, we assume that the twin disks pass through the vertices of τ and that the angle at s^* measures exactly $\pi/3$. Now it is not hard to show that R is maximized exactly when T is equilateral: if we move s^* away from the midpoint of the arc covered by the twin disks, keeping the twin disks passing through the vertices of T, the radius R decreases, until s^* reaches one of the other

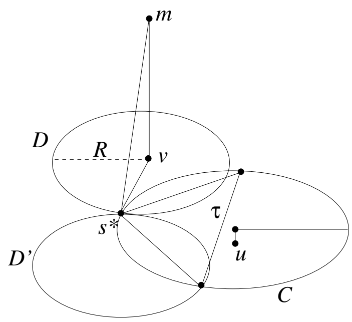

Figure 12: Illustration for the proof of Lemma 8.5.

vertices of τ and $R = \mathrm{radius}(C)$. Since the worst-case configuration is equilateral T, we can conclude that $R \leq \sqrt{3}\mathrm{radius}(C)$.

Now we can bound all these radii. Let u denote the restricted Voronoi vertex dual to T. Since u is the minimum circumradius of τ, $\mathrm{radius}(C) \leq d(u, s^*)$. By Corollary 8.1, $d(u, s^*) = O(\epsilon^{\gamma/8})$, which implies that $R = O(\mathrm{radius}(C)) = O(\epsilon^{\gamma/8})$.

Now to find the angle between the normal to τ and the normal to S_α at s^*, we consider one of the twin tangent balls B at s^*. Let m denote the center of B and let v denote the center of D. The segment s^*m is normal to S_α at s^* and the segment mv is normal to T, so the angle $\angle(n_\tau, n_{s^*})$ we would like to bound is $\angle s^*mv$. Observe that $\angle s^*mv = \pi/2$, $d(s^*, m) = 1/10$, and $d(s^*, v) = R = O(\epsilon^{\gamma/8})$. Hence $\angle(n_\tau, n_{s^*}) = \angle s^*mv \leq \arcsin(O(\epsilon^{\gamma/8})) = O(\epsilon^{\gamma/8})$.

Take any vertex p^* of τ other than s^*. By Corollary 8.1, $d(u, s^*) = O(\epsilon^{\gamma/8})$ and $d(u, p^*) = O(\epsilon^{\gamma/8})$. Thus $d(s^*, p^*) \leq d(u, s^*) + d(u, p^*) = O(\epsilon^{\gamma/8})$. Then

$$\begin{aligned} d(\tilde{s^*}, \tilde{p^*}) &\leq d(s^*, p^*) + d(s^*, \tilde{s^*}) + d(p^*, \tilde{p^*}) \\ &= O(\epsilon^{\gamma/8}) + 2\Delta \\ &= O(\epsilon^{\gamma/8}) + O(\epsilon^{\gamma/2}) \\ &= O(\epsilon^{\gamma/8}). \end{aligned}$$

And by Lemma 2.2(ii), we have $\angle(n_{s^*}, n_{p^*}) \leq 3d(\tilde{s^*}, \tilde{p^*}) = O(\epsilon^{\gamma/8})$. Hence $\angle(n_T, n_{p^*}) \leq \angle(n_T, n_{s^*}) + \angle(n_{s^*}, n_{p^*}) = O(\epsilon^{\gamma/8})$. ◻

To prove Lemma 8.6, we need Lemma 8.5 and two more lemmas from paper [1] as stated below.

Lemma 9.5 *For any two points p and q on S with $d(p, q) \leq \lambda f(p)$, $\angle(pq, n_p) \geq \pi/2 - \arcsin(\lambda/2)$.*

Lemma 9.6 *If a ball B intersects surface S in more than one connected component, then B contains a point of the medial axis of S.*

With these lemmas, by using almost the same arguments as Theorem 2 in [1], we can prove that the closed ball property: the closure of each k-dimensional face, $1 \leq k \leq 3$, of the Voronoi diagram of C^* intersects intersect S in either the empty set *or* in a closed $(k-1)$-dimensional topological ball. With this property, by the theorem (Theorem 4.3) of Edelsbrunner and Shah [13], we immediately have Lemma 8.6.

Proof of Lemma 8.6

To prove the closed ball property, we will consider the intersection of a Voronoi edge, a Voronoi face, and a Voronoi cell with the surface S respectively.

Let p^* be a selected center point in C^*, and let V_{p^*} be its Voronoi cell. Let the direction of n_{p^*} be vertical.

We begin by showing that in the vicinity of p^*, the surface S is nearly horizontal. Corollary 8.1 shows that $V_{p^*}|_S$ is small, fitting inside a ball B centered at p^* of radius $O(\epsilon^{\gamma/8})$. Since such a small ball cannot intersect the medial axis, Lemma 9.6 implies that $S \cap B$ is connected. Lemma 2.2(ii) shows that the normal to $S \cap B$ makes an angle $O(\epsilon^{\gamma/8})$ radians with the vertical.

Consider an edge e of V_{p^*}. If e intersects S, then since e is normal to the dual Delaunay triangle T, Lemma 8.5 implies that e is within $O(\epsilon^{\gamma/8})$ radians from the normal to S at p^*. So e can intersect S only once within ball B.

Next consider a face f of V_{p^*}. The face f is contained in its supporting plane H, the perpendicular bisector of p^* and another center point s^*, where p^*s^* is an edge of a restricted triangle T. The plane H must contain a vector h parallel to the normal of T, so again Lemma 8.5 shows that the angle $\angle(h, n_{p^*})$ is at most $O(\epsilon^{\gamma/8})$ radians. Each component of $f \cap S$ is an arc of a curve, with endpoints on the edges of f. Assume for the sake of contradiction that there are at least two such connected components, and consider any line segment connecting a point on one component with a point on another. Since $V_{p^*}|_S$ is small, Lemma 9.5 implies that each of these line segment is *nearly horizontal*, specifically within angle $O(\epsilon^{\gamma/8})$ radians with the horizontal. Hence we can sort the arcs of $f \cap S$ from left to right across f, as shown in Figure 9(a). Let q_1q_2 be a line segment connecting the right endpoint of one arc with the left endpoint of the next arc. The segment q_1q_2 is nearly horizontal so it must leave f as it crosses the nearly vertical edge of V_{p^*} at q_1 and reenter f at q_2, a contradiction to the fact that f is convex.

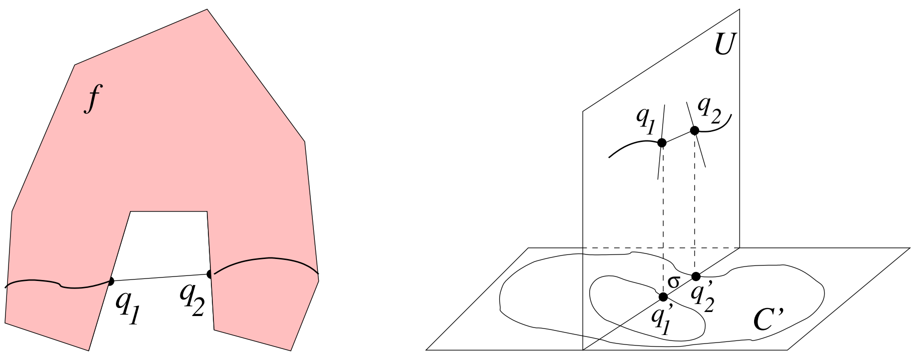

Figure 13: Two impossible situations: (a) Face f of V_{p^*} intersects S in two arcs. (b) $V_{p^*} \cap S$ is a disk with a hole.

Finally consider V_{p^*} itself. Let C be $V_{p^*}|_S$. Now C cannot have a handle because S is nearly horizontal everywhere within ball B. We assume, however, that C is not a topological disk, again aiming for a contradiction. If C has no handles and is not a topological disk, then either it has more than one connected component or it is a topological disk with holes.

Consider the projection C' of C onto a horizontal plane. Since each pair of points in C are connected by a nearly horizontal segment, this projection is one-to-one, and C' is a planar shape homeomorphic to C. If C' has more than one connected component, let σ be the shortest segment connecting two different components, and let q_1', q_2' be that endpoints of σ. Otherwise, C' has a hole. Select some point x in a hole, let σ' be the line segment connecting x to the nearest point q_1' in C'. The line supporting σ' intersects C' again on the other side of x in a point q_2'. We let σ be the segment $q_1'q_2'$. Please refer to Figure 9(b).

Segment σ is perpendicular to the boundary of C' at q_1'. Let U be the vertical plane through σ, and let q_1, q_2 be the points in $U \cap C$ whose vertical projections are q_1', q_2' respectively. The point q_1 is either a vertex of C or an edge point, and thus lies in either an edge or facet of V_{p^*} respectively.

Consider the case in which q_1 lies in a facet f of V_{p^*}. Facet f is nearly vertical, and since U is perpendicular to C' at q_1', it is nearly perpendicular to C at q_1, and f intersects U in a nearly vertical line at q_1. In the other case, q_1 is contained in a nearly vertical edge e of V_{p^*}. We consider the plane V containing both e and the horizontal line perpendicular to segment q_1q_2 at q_1. V meets U in a nearly vertical line. Note that we had to choose U carefully, since two nearly vertical planes which are not nearly perpendicular might meet in a nearly horizontal line.

In either case, examining the situation in U, we find that in the neighborhood of q_1, the interior of q_1q_2 is separated from the interior of $V_{p^*} \cap U$ by a nearly vertical line. However, both q_1 and q_2 belong to $V_{p^*} \cap U$, contradicting to the fact that V_{p^*} is convex. ◻

Proof of Lemma 8.9

Suppose p^* is on the surface S_α, where $\alpha \leq \Delta$. Let B be the Voronoi ball centered at v. Let $M_1 = B(m_1, 1/10), M_2 = B(m_2, 1/10)$ be the twin empty balls tangent to S_α at p^* in Lemma 8.4. Suppose v and m_1 lie on the same side of S_α. Refer to Figure 14.

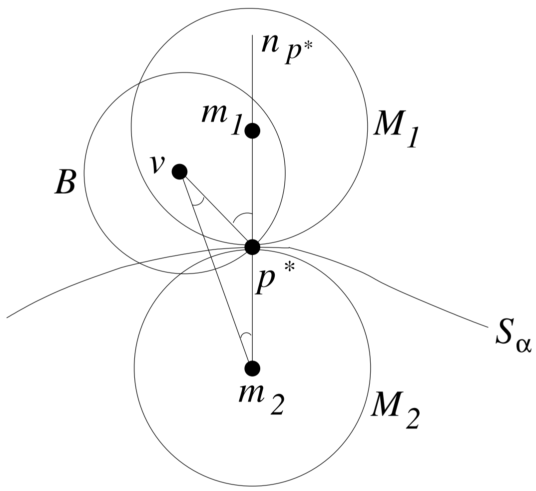

Figure 14: Illustration for the proof of Lemma 8.9.

As m_2 and v lie on opposite sides of S_α, line segment m_2v must intersect S_α at a point, say q.

As M_1 and M_2 are empty, m_1 and m_2 belong to V_{p^*}. So $m_2 v \subset V_{p^*}$, which implies $q \in V_{p^*}$. By Corollary 8.1, $d(q, p^*) = O(\epsilon^{\gamma/8})$. It is clear that $d(q, p^*) \geq d(v, p^*) \sin \angle v \geq \epsilon^{\gamma/16} \sin \angle v$. Thus $\angle v = O(\epsilon^{\gamma/8})/\epsilon^{\gamma/16} = O(\epsilon^{\gamma/16})$. On the other hand, $d(q, p^*) \geq d(v, p^*) \sin \angle m_2 = (\sin \angle m_2)/10$. Thus $\angle m_2 = O(\epsilon^{\gamma/8})$. Hence $\angle(p^* v, n_{p^*}) = \angle v p^* m_1 \leq \angle m_2 + \angle v = O(\epsilon^{\gamma/16}) + O(\epsilon^{\gamma/8}) = O(\epsilon^{\gamma/16})$. ◻

Proof of Lemma 8.11

Suppose $p^* \in S_\alpha$. Consider the twin empty balls M_1 and M_2 touching S_α at p^* of radius $1/10$. Let D be the ball with the circumcircle of T as a diameter; refer to Figure 15(a). The radius r of D is equal to

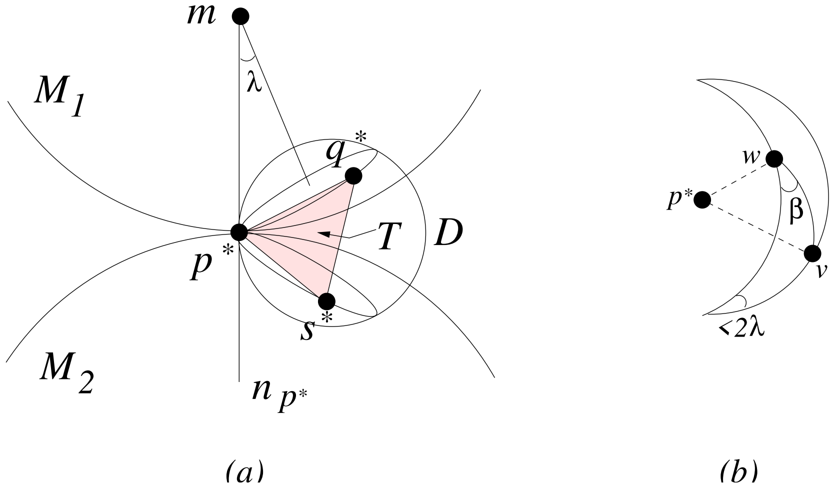

Figure 15: Normal of a small triangle T and the normal to S_α at the vertex p^* of T with the largest vertex angle.

the radius of the circumcircle of T. Denote the circles of the intersection of D with M_1 and M_2 as C_1 and C_2 respectively. The normal to S at p passes through m, the center of M_1. This normal makes an angle less than λ with the normals to the planes of C_1 and C_2, where

$$\begin{aligned} \lambda &\leq \arcsin(r/d(p^*, m)) \\ &= \arcsin(10r) \\ &= O(\epsilon^{\gamma/16}) \end{aligned}$$

by Lemma 8.10. This angle bound also applies to the plane of C_2, which implies that the planes of C_1 and C_2 make a wedge, say G, with an acute angle no more than 2λ.

The other two vertices q^* and s^* of T cannot lie inside M_1 and M_2. This implies that T lies completely in the wedge G. Consider a cone at p^* inside the wedge G formed by the three planes: the plane π_T supporting T, the plane π_1 supporting C_1, and the plane π_2 supporting C_2. A unit sphere centered around p^* intersects the cone in a spherical triangle uvw, where u, v and w are the intersection points of the lines $\pi_1 \cap \pi_2$, $\pi_T \cap \pi_1$ and $\pi_T \cap \pi_2$ with the unit sphere respectively. See Figure 15(b). Without the loss of generality, assume that the angle $\angle uvw \leq \angle uwv$. We have the following facts. The arc length of wv, denoted by $|wv|$, is at least $\pi/3$ since p^* subtends the largest angle in T and T lies completely in the wedge G. The spherical angle $\angle vuw$ is less than or equal to 2λ. We are interested in the spherical angle $\beta = \angle uvw$ which is also the acute dihedral angle between the planes of T and C_1. By standard *sine laws* in spherical geometry, we have $\sin\beta = \sin|uw| \frac{\sin\angle vuw}{\sin|wv|} \leq \sin\sin|uw| \frac{\sin(2\lambda)}{\sin|wv|}$. If $\pi/3 \leq |wv| \leq 2\pi/3$, we have $\sin\pi/3 \geq \sqrt{3}/2$ and hence $\beta \leq \arcsin\frac{2}{\sqrt{3}}\sin(2\lambda)$. For the range $2\pi/3 < |wv| \leq \pi$, we use the fact that $|uw| + |wv| \leq \pi$ since $\angle vuw \leq 2\alpha < \pi/2$ for sufficiently small ϵ. So, in this case $\frac{\sin|uw|}{\sin|wv|} < 1$. Thus $\beta \leq \arcsin\frac{2}{\sqrt{3}}\sin(2\lambda) = O(\epsilon^{\gamma/16})$.

Hence the angle between the normal of T and n_{p^*} is an acute angle at most $\lambda + \beta = O(\epsilon^{\gamma/16})$. $\square$

Proof of Lemma 8.18

We know that μ is well-defined and continuous on U, and its inverse is also continuous, so it only remains to show that it is one-to-one. First, we prove that if U is in one triangle T, μ is one-to-one. For a point $q \in T$, the vector from q to $\mu(q)$ is normal to the surface at $\mu(q)$. If there was some $y \neq q$ on T with $\mu(y) = \mu(q)$, then q, $\mu(q)$ and y would all be collinear and T itself would have to contain line segment qy, contradicting Lemma 8.12(ii), which says that the normal of T is nearly parallel to the normal of S at $\mu(q)$.

Now, we consider the case in which U is contained in more than one triangle. Let q and y be two points U such that $\mu(q) = \mu(y) = x$, and let p^* be a common vertex of the triangles that contain U. Since μ is one-to-one in one triangle, q and y must lie in the two distinct triangles T_q and T_y respectively. Let l be the normal line to S at x. Note that $l \parallel n_x$. l pieces the patch U at least twice; if y and q are not adjacent intersections along l, redefine q so that this is true. Now consider the orientation of the patch U according to the estimated normal direction ℓ_p^* at p^*. Either l passes from inside to outside and back to inside when crossing y and q, or from outside to inside and back to outside.

The acute angles between the triangle normals of T_q, T_y and n_x are $O(\epsilon^{\gamma/16})$ by Lemma 8.12(ii), that is, the triangles are stabbed nearly perpendicularly by n_x. But since the orientation of U is opposite at the two intersections, the angle between the two *oriented* triangle normals is greater than zero, meaning that T_q and T_y must meet at p^* at an acute angle. This would contradict Lemma 8.14, which is that T_q and T_y meet at p^* at an obtuse angle. Hence there are no two points y, q in U with $\mu(q) = \mu(y)$. $\square$

Proof of Lemma 8.19

Let $S' \subset S$ be $\mu(N)$. We first show that (N, μ) is a *covering space* of S'. Informally, (N, μ) is a covering space for S' if function μ maps N smoothly onto S', with no folds or other singularities; see Massey [22], Chapter 5. Showing that (N, μ) is a covering space is weaker than showing that μ defines a homeomorphism, since, for instance, it does not preclude several connected components of N mapping onto the same component of S', or more interesting behavior, such as torus wrapping twice around another torus wrapping twice around another torus to form a *double covering*.

Formally, the (N, μ) is a covering space of S' if, for every $x \in S'$, there is a path-connected *elementary neighborhood* V_x around x such that each path-connected component of $\mu^{-1}(V_x)$ is mapped homeomorphically onto V_x by μ.

To construct such an elementary neighborhood, note that the set of points $|\mu^{-1}(x)|$ corresponding to a point $x \in S'$ is non-zero and finite, since μ is one-to-one on each triangle of N and there are only a finite number of triangles. For each point $q \in \mu^{-1}(x)$, we choose an open neighborhood U_q of around

q, homeomorphic to a disk and small enough so that U_q is contained only in triangles that contain q. Refer to Figure 16.

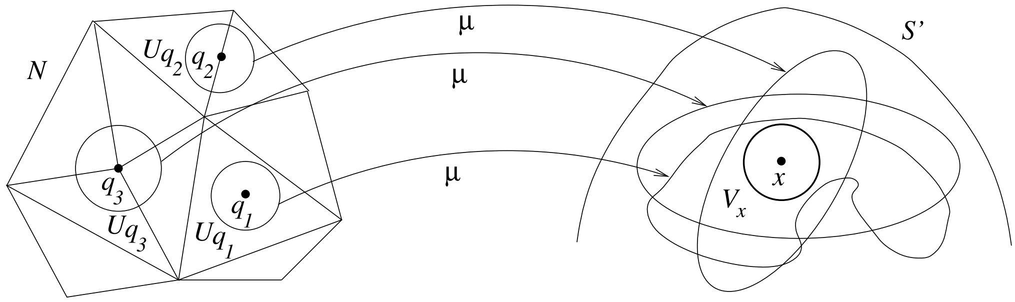

Figure 16: Illustration for the proof of Lemma 8.19.

As m_2 and v lie on opposite sides of S_α,

We claim that μ maps each U_q homeomorphically onto $\mu(U_q)$. This is because it is continuous, it is onto $\mu(U_q)$ by definition, and, since any two points x and y in U_q are in adjacent triangles, it is one-to-one by Lemma 8.18.

Let $U'(x) = \cap_{q \in \mu^{-1}(x)} \mu(U_q)$, the intersection of the maps of each of the U_q. $U'(x)$ is the intersection of a finite number of open neighborhoods, each containing x, so we can find an open disk V_x around x. V_x is path connected, and each component of $\mu^{-1}(V_x)$ is a subset of some U_q and hence is mapped homeomorphically onto V_x by μ. Thus (N, μ) is a covering space for S'.

We show that μ defines a homeomorphism between N and S'. Since N is onto S' by definition, we need only show that μ is ono-to-one. Consider one connected component G of S'. A theorem of algebraic topology (see eg. Massey [22], Chapter 5, Lemma 3.4) says that when (N, μ) is a covering space of S', the sets $\mu^{-1}(x)$ for all $x \in G$ have the same cardinality. We now use Corollary 8.2(ii), that μ is one-to-one from N to $\mu(p^*)$ for every $p^* \in C^*$. From the algorithm, it is clear that each connected component of S contains a $\mu(p^*)$ for some center point $p^* \in C^*$. So it must be the case that $\mu : N \to S'$ is everywhere one-to-one, and N and S' are homeomorphic.

Finally, we show that $S' = S$. Since N is closed and compact, S' must be as well. So S' cannot include part of a connected component of S, and hence S' must consist of a subset of the connected components of S. Since every connected component of S contains a $\mu(p^*)$ for some center point $p^* \in C^*$, all components of S belong to S'. Hence $S' = S$, and N and S are homeomorphic. ◻