Curve Reconstruction from Noisy Samples

Siu-Wing Cheng* Stefan Funke†‡ Mordecai Golin*
Piyush Kumar§ Sheung-Hung Poon* Edgar Ramos¶ ‡

November 4, 2003

Abstract

We present an algorithm to reconstruct a collection of disjoint smooth closed curves from noisy samples. Our noise model assumes that the samples are obtained by first drawing points on the curves according to a locally uniform distribution followed by a uniform perturbation in the normal directions. Our reconstruction is faithful with probability approaching 1 as the sampling density increases. We expect that our approach can lead to provable algorithms under less restrictive noise models and for handling non-smooth features.

1 Introduction

The combinatorial curve reconstruction problem has been extensively studied recently by computational geometers. The input consists of sample points on a collection of unknown disjoint smooth closed curves denoted by F. The problem calls for computing a set of polygonal curves that are provably *faithful*. That is, as the sampling density increases, the polygonal curves should converge to F.

Amenta et al. [3] obtained the first results in this problem. They proposed a *2D crust* algorithm whose output is provably faithful if the input satisfies the *ϵ-sampling* condition for any $\epsilon < 0.252$. For each point x on F, the *local feature size* $f(x)$ at x is defined as the distance from x to the medial axis of F. For $0 < \epsilon < 1$, a set S of samples is an ϵ-sampling of F if for any point $x \in F$, there exists a sample $s \in S$ such that $\|s - x\| \leq \epsilon \cdot f(x)$ [3]. The algorithm by Amenta et al. invokes the computation of a Voronoi diagram or Delaunay triangulation twice. Gold and Snoeyink [11] simplified the algorithm and invokes the computation of Voronoi

*Department of Computer Science, HKUST, Hong Kong. Email:{scheng,golin,hung}@cs.ust.hk. Research of S.-W. Cheng and S.-H. Poon are partly supported by Research Grant Council, Hong Kong, China (project no. HKUST 6190/02E). Research of M. Golin is partly supported by Research Grant Council, Hong Kong, China (project no. HKUST 6082/01E and HKUST 6206/02E).

†MPI, Saarbrücken, Germany. Email: funke@mpi-sb.mpg.de

‡Partly supported by the IST Programme of the EU as a Shared-cost RTD (FET Open) Project under Contract No IST-2000-26473 (ECG - Effective Computational Geometry for Curves and Surfaces).

§Department of Computer Science, SUNY, USA. Email:piyush@ams.sunysb.edu. Research supported by National Science Foundation grant CCR-0098172.

¶Department of Computer Science, UIUC, USA. Email:eramosn@cs.uiuc.edu

diagram or Delaunay triangulation only once. Later, Dey and Kumar [4] proposed a *NN-crust* algorithm for this problem. Since we will use the NN-crust algorithm, we briefly describe it. For each sample s in S, connect s to its nearest neighbor in S. Afterwards, if a sample s is incident on only one edge e, connect s to the closest sample among all samples u such that su makes an obtuse angle with e. The output curve is faithful for any $\epsilon \leq 1/3$ [4].

Dey, Mehlhorn, and Ramos [5] proposed a *conservative-crust* algorithm to handle curves with endpoints. Funke and Ramos [9] proposed an algorithm to handle curves that may have sharp corners and endpoints. Dey and Wenger [6, 7] also described algorithms and implementation for handling sharp corners. Giesen [10] discovered that the traveling salesperson tour through the samples is a faithful reconstruction, but this approach cannot handle more than one curve. Althaus and Mehlhorn [2] showed that such a traveling salesperson tour can be constructed in polynomial time.

Noise often arises in collecting the input samples. For example, when the input samples are obtained from 2D images by scanning. The noisy samples are typically classified into two types. The first type are samples that cluster around F but they generally do not lie on F. The second type are outliers that lie relatively far from F. No combinatorial algorithm is known so far that can compute a faithful reconstruction in the presence of noise. In this paper, we propose a method that can handle noise of the first type for a set of disjoint smooth closed curves. We assume that the input does not contain outliers. Proving a deterministic theorem seems difficult as arbitrary noisy samples can collaborate to form patterns to fool any reconstruction algorithm. Instead, we assume a particular model of noise distribution and prove that our reconstruction is faithful with probability approaching 1 as the number of samples increases. For simplicity and notational convenience, we assume throughout this paper that $\min_{x \in F} f(x) = 1$ and F consists of a single smooth closed curve, although our algorithm works when F contains more than one curve.

In our model, a sample is generated by drawing a point from F followed by randomly perturbing the point in the normal direction. Let $L = \int_F \frac{1}{f(x)} dx$. The drawing of points from F follows the probability density function $\frac{1}{L \cdot f(x)}$. That is, the probability of drawing a point from a curve segment η is equal to $\int_\eta \frac{1}{f(x)} dx$ divided by L. A point p drawn from F is then perturbed in the normal direction. The perturbation is uniformly distributed within an interval that has p as the midpoint, width 2δ, and aligns with the normal direction at p. The distribution of each sample is independently identical. δ is the noise amplitude and we assume that $\delta \leq 1/(9\rho^2)$ where $\rho \geq 4$ is a constant chosen a priori by our algorithm. We assume throughout this paper that $\delta > 0$. We emphasize that the value of δ is unknown to our algorithm. Although the perturbation along the normal direction is restrictive, it isolates the effect of noise from the distribution of samples on F. This facilities an initial study of curve reconstruction in the presence of noise.

We prove that our algorithm returns a reconstruction which is faithful with probability at least $1 - O(n^{-\Omega(\frac{\ln^\omega n}{f_{\max}} - 1)})$, where n is the number of input samples, ω is an arbitrary positive constant, and $f_{\max} = \max_{x \in F} f(x)$. Our algorithm works for noisy samples from a collection of disjoint smooth closed curves. The novelty of our algorithm is a method to cluster samples so that each cluster comes from a relatively flat portion of F. This allows us to estimate points that lie close to F. We believe that this clustering approach will also be useful for less restrictive noise models and recognizing non-smooth features. We also expect that this clustering approach can be generalized to 3D for surface reconstruction problems.

The rest of the paper is organized as follows. Section 2 describes our algorithm. Section 3 introduces two decompositions of the space around F which is the main tool in our probabilistic argument. Sections 4 and 5 prove that our reconstruction is faithful with probability approaching 1. Section 6 discusses extension to handling non-smooth features.

2 Algorithm

We first highlight the key ideas. Our algorithm works by growing a disk neighborhood around each sample p until the samples inside the disk fit in a strip whose width is small relative to the radius of the disk. The final disk is the *coarse neighborhood* of p denoted by *coarse*(p). *coarse*(p) provides a first estimate of the curve locally and of its normal. A better estimation is possible. We shrink *coarse*(p) by a certain factor. We take a slab bounded by two parallel tangent lines of the shrunken *coarse*(p). The slab is the *refined neighborhood* of p denoted by *refined*(p). We rotate *refined*(p) around p to minimize the spread of the samples in *refined*(p) along the direction of *refined*(p). The final orientation of *refined*(p) provides a good normal estimation and it also allows us to estimate a *center point* close to F in place of p. Next, we

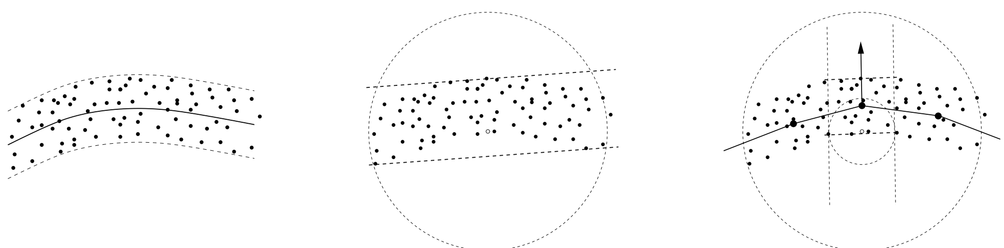

Figure 1: On the left, a smooth curve segment with a noise cloud. In the middle, a sufficiently large neighborhood identifies a strip with relatively large aspect ratio, which can provide preliminary point and normal estimates. On the right, concentrating on smaller neighborhoods, a better estimate of point and normal is possible.

decimate the center points as follows. We scan the center points in decreasing order of the widths of their corresponding refined neighborhoods. When we add the current center point p^* to the decimated set, we delete the other center points that are too close to p^*. Finally, we can run any reconstruction algorithm that is correct for a noise free sampling on the remaining center points. For example, the NN-Crust algorithm by Dey and Kumar [4].

We provide the details of the algorithm in the following. Let n be the total number of input samples. Let $\omega > 0$ and $\rho \geq 4$ be two predefined constants.

Point Estimation: For each sample s, we estimate a point as follows.

Coarse neighborhood: Let D be the disk that is centered at s and contains $\ln^{1+\omega} n$ samples. Let *initial*(s) be the disk centered at s with radius $\sqrt{\text{radius}(D)}$. We initialize *coarse*(s) = *initial*(s) and compute an infinite strip *strip*(s) of minimum width that contains all samples inside *coarse*(s). We grow *coarse*(s) and maintain *strip*(s) until $\frac{\text{radius}(\textit{coarse}(s))}{\text{width}(\textit{strip}(s))} \geq \rho$. The final disk *coarse*(s) is the *coarse neighborhood* of s.

Refined neighborhood: Let N_s be a direction perpendicular to the long side of *strip*(s). The *refined neighborhood refined*(s) is the slab that contains s in the middle, parallel to N_s, and has width equal to $\min\{\sqrt{\text{radius}(\textit{initial}(s))}, \text{radius}(\textit{coarse}(s))/3\}$. We enclose the samples in *refined*(s) by two parallel lines that are orthogonal to N_s. These two lines form a rectangle *rectangle*(s) with the boundary lines of *refined*(s). We rotate *refined*(s) around s in the clockwise and anti-clockwise directions and maintain *rectangle*(s) . The range of the rotation is $[0, \pi/10]$. Within this range, we position *refined*(s) such that the height of *rectangle*(s) in the direction N_s is minimized. We return the center point s^* of the final *rectangle*(s).

Pruning: We sort the center points s^* in decreasing order of width(*refined*(s)). Then we scan the sorted list and select a subset of center points: when we select the current center point s^*, we delete all center points u^* from the sorted list such that $\|s^* - u^*\| \leq \text{width}(\textit{refined}(s))^{1/3}$.

Output: We run the NN-crust algorithm on the selected center points and return the output curve.

3 Decompositions

For each point $x \in \mathbb{R}^2$ that does not lie on the medial axis of F, we use $\tilde{x}$ to denote the point on F closest to x. That is, $\tilde{x}$ is the projection of x onto F. (We are not interested in points on the medial axis.)

We call the bounded region enclosed by F the *inside* of F and the unbounded region the *outside* of F. For $0 < \alpha \leq \delta$, F_α^+ (resp. F_α^-) is the curve that passes through the points q inside (resp. outside) F such that $\|q - \tilde{q}\| = \alpha$. We use F_α to mean F_α^+ or F_α^- when it is unimportant to distinguish between inside and outside. The *normal segment* at a point $p \in F$ is the line segment consisting of points q on the normal of F at p such that $\|p - q\| \leq \delta$. Given two points x and y on F, we use $F(x, y)$ to denote the curved segment traversed from x to y in clockwise direction. We use $|F(x, y)|$ to denote the length of $F(x, y)$.

We will use two types of decompositions, *β-partition* and *β-grid*. Let $0 < \beta < 1$ be a parameter. We identify a set of *cut-points* on F as follows. We pick an arbitrary point c_0 on F as the first cut-point. Then for $i \geq 1$, we find the point c_i such that c_i lies in the interior of $F(c_{i-1}, c_0)$, $|F(c_{i-1}, c_i)| = \beta^2 f(c_{i-1})$, and $|F(c_i, c_0)| \geq \beta^2 f(c_i)$. If c_i exists, it is the next cut-point and we continue. Otherwise, we have computed all the cut-points and we stop. The β-partition is the arrangement of the normal segments at the cut-points, F_δ^+, and F_δ^-. Figure 2 shows an example. We call each face of the β-partition a β-*slab*. The β-partition consists of a row of slabs stabbed by F.

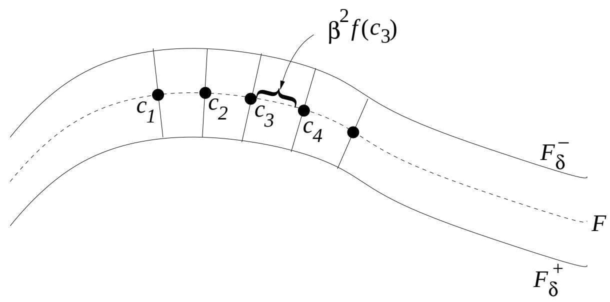

Figure 2: β-partition.

The cut-points for a β-grid are picked differently. We pick an arbitrary point c_0 on F as the first cut-point. Then for $i \geq 1$, we find the point c_i such that c_i lies in the interior of $F(c_{i-1}, c_0)$, $|F(c_{i-1}, c_i)| = \beta f(c_{i-1})$, and $|F(c_i, c_0)| \geq \beta f(c_i)$. If c_i exists, it is the next cut-point and we continue. Otherwise, we have computed all the cut-points and we stop. The β-grid is the arrangement of the following:

- The normal segments at the cut-points.
- F, F_δ^+, and F_δ^-.
- F_α^+ and F_α^- where $\alpha = i\beta\delta$ and i is an integer between 1 and $\lfloor 1/\beta \rfloor - 1$.

The β-grid has a grid structure. Figure 3 shows an example. We call each face of the β-grid a β-*cell*. There are $O(1/\beta)$ rows of cells "parallel to" F.

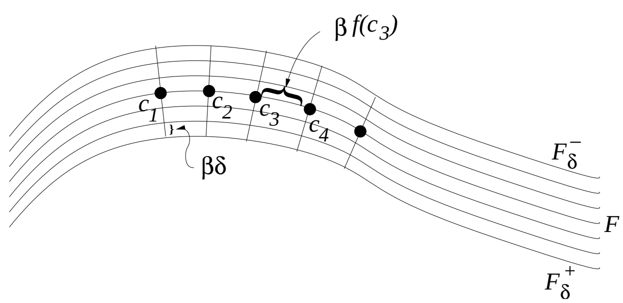

Figure 3: β-grid.

In Section 3.1, we prove several properties of F_α for any α. These properties will be used in Section 3.2 to bound the diameter of a β-cell. These properties will also be useful later in the paper. In Section 3.3, we analyze the probabilities of a β-slab and a β-cell containing certain numbers of samples. These probabilities are essential for the probabilistic analysis later.

3.1 Properties of F_α

Lemma 3.1 *Any point p on F_α has two tangent disks with radii $f(\tilde{p}) - \alpha$ whose interior do not intersect F_α.*

Proof. Let M_α be the medial disk of F_α touching a point $p \in F_\alpha$. By the definition of F_α, there is a medial disk M of F touching $\tilde{p}$ such that M and M_α have the same center and $\text{radius}(M_\alpha) = \text{radius}(M) - \alpha \geq f(\tilde{p}) - \alpha$. Let D be a disk of radius $f(\tilde{p}) - \alpha$ that touches F_α at p. If F_α intersects the interior of D, the medial axis of F_α intersects the interior of D. So $\text{radius}(M_\alpha) < \text{radius}(D) = f(\tilde{p}) - \alpha$, contradiction. ◻

For each point p on F_α, define *cocone*(p, θ) as the double cone that has apex p and angle θ such that the normal of F_α at p is the symmetry axis of the double cone that lies outside the double cone. The next lemma shows that F_α lies inside *cocone*(p, θ) for a small θ in a small neighborhood of p.

Lemma 3.2 *Let p be a point on F_α. Let D be a disk centered at p with radius at most $2(1-\alpha)f(\tilde{p})$.*

(i) For any point $q \in F_\alpha \cap D$, the distance of q from the tangent at p is at most $\frac{\|p-q\|^2}{2(1-\alpha)f(\tilde{p})}$.

(ii) $F_\alpha \cap D \subseteq$ cocone$(p, 2\sin^{-1}\frac{\text{radius}(D)}{2(1-\alpha)f(\tilde{p})})$.

Proof. Assume that the tangent at p is horizontal. Consider (i). Refer to Figure 4(a). Let B be the tangent disk at p that lies above p and has center x and radius $(1-\alpha)f(\tilde{p})$. Let C be the circle centered at p with radius $\|p-q\|$. Since $\|p-q\| < 2(1-\alpha)f(\tilde{p})$, C crosses B. Let r be a point in $C \cap \partial B$. Let d be the distance of r from the tangent at p. By Lemma 3.1, d bounds the distance from q to the tangent at p. Observe that $\|p-q\| = \|p-r\| = 2(1-\alpha)f(\tilde{p})\sin(\angle pxr/2)$ and $d = \|p-r\| \cdot \sin(\angle pxr/2)$. Thus, $d = 2(1-\alpha)f(\tilde{p})\sin^2(\angle pxr/2) = \frac{\|p-q\|^2}{2(1-\alpha)f(\tilde{p})}$.

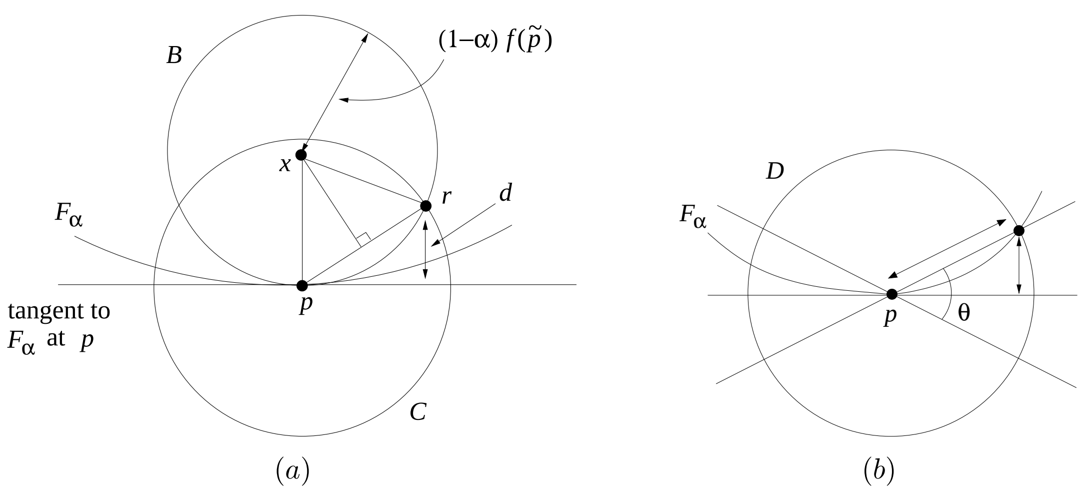

Figure 4: Illustration for Lemma 3.2.

Consider (ii). Refer to Figure 4(b). By (i), the distance between $F_\alpha \cap D$ and the tangent at p is bounded by $\frac{\text{radius}(D)^2}{2(1-\alpha)f(\tilde{p})}$. Let θ be the smallest angle such that *cocone*(p, θ) contains $F_\alpha \cap D$. Then $\sin\frac{\theta}{2} \leq \frac{\text{radius}(D)^2}{2(1-\alpha)f(\tilde{p})} \cdot \frac{1}{\text{radius}(D)} = \frac{\text{radius}(D)}{2(1-\alpha)f(\tilde{p})}$. ◻

The next lemma shows that the normal deviation is very small in a small neighborhood of any point in F_α.

Lemma 3.3 *Let p be a point on F_α. Let D be a disk centered at p with radius at most $\frac{(1-\alpha)f(\tilde{p})}{4}$. For any point $u \in F_\alpha \cap D$, the acute angle between the normals at p and u is at most $2\sin^{-1}\frac{\|p-u\|}{(1-\alpha)f(\tilde{p})} \leq 2\sin^{-1}\frac{\text{radius}(D)}{(1-\alpha)f(\tilde{p})}$.*

Proof. Take any point u on $F_\alpha \cap D$. Let ℓ be the tangent to F_α at u. Let ℓ' be the line that is perpendicular to ℓ and passes through u. Let C be the circle centered at p with radius $\|p-u\|$. Let A and B be the two tangent circles at p with radius $\frac{(1-\alpha)f(\tilde{p})}{2}$. Let x be the center of A. Without loss of generality, we assume that the tangent to F_α at p is horizontal, A is below B, u lies to the left of p, and the slope of ℓ is positive or infinite. (We ignore the case where the slope of ℓ is zero as there is nothing to prove then.) It follows that the slope of ℓ' is zero or negative. Refer to Figure 5.

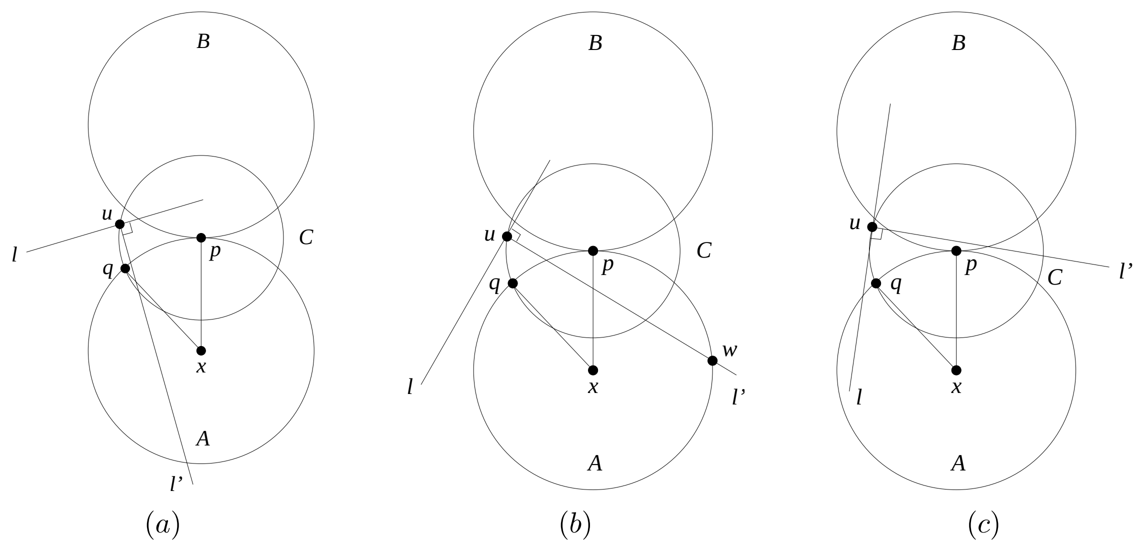

Figure 5: Illustration for Lemma 3.3.

By Lemma 3.1, u lies outside A and B. Let q be the intersection point between C and A on the left of p. Since $\|p-q\| = \|p-u\| \leq \frac{(1-\alpha)f(\tilde{p})}{4} = \text{radius}(A)/2$, q lies above x. Also, $\angle pxq = 2\sin^{-1}\frac{\|p-u\|}{(1-\alpha)f(\tilde{p})}$.

Suppose that ℓ' does not lie above x, see Figure 5(a). Since u lies above the support line of qx, the angle between ℓ' and the vertical is less than or equal to $\angle pxq = 2\sin^{-1}\frac{\|p-u\|}{(1-\alpha)f(\tilde{p})}$.

Suppose that ℓ' lies above x but not above p, see Figure 5(b). We show that this case is impossible. Let w the intersection point between A and ℓ' on the right of p. Note that p lies between u and w and $\angle upw > \pi/2$. If we grow a disk that lies below l and remains tangent to l at u, the disk will hit F_α at some point different from u when the disk passes through p or earlier. It follows that there is a medial disk M_u of F_α that touches u and lies below l. Observe that the center of M_u lies on the half of ℓ' on the right of u. Furthermore, the center of M_u

lies on the line segment uw; otherwise, since $\angle upw > \pi/2$, M_u would contain p, contradiction. Thus, the distance from $\tilde{p}$ to the center of M_u is less than $\max\{\|p-u\|, \|p-w\|\} + \|p-\tilde{p}\| \leq 2 \cdot \text{radius}(A) + \alpha = (1-\alpha)f(\tilde{p}) + \alpha \leq f(\tilde{p})$. But the center of M_u is also a point on the medial axis of F which implies that $f(\tilde{p}) < f(\tilde{p})$, contradiction.

The remaining case is that ℓ' lies above p, see Figure 5(c). Since u lies outside B and the slope of ℓ' is zero or negative, ℓ' lies between p and the center of B. The situation is similar to the previous case where ℓ' lies between p and x. So a similar argument shows that this case is also impossible. ◻

3.2 Diameter of a β-cell

In this section, we prove an upper bound on the diameter of a β-cell. First, we need a utility lemma.

Lemma 3.4 *Assume that $\beta \leq 1/4$. Let p and q be two points on F_α such that $|F(\tilde{p},\tilde{q})| \leq 2\beta f(\tilde{p})$. Then $\|p-q\| \leq \|\tilde{p}-\tilde{q}\| + 5\beta\delta$.*

Proof. Refer to Figure 6. Let r be the point $q - \tilde{q} + \tilde{p}$. Without loss of generality, assume that $\angle\tilde{p}pr \leq \angle\tilde{p}rp$. Lemma 3.3 implies that $\angle\tilde{p}pr \leq 2\sin^{-1}2\beta$. Therefore, $\angle\tilde{p}rp \geq \pi/2 -$

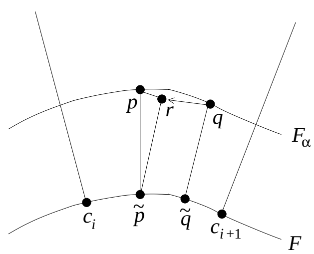

Figure 6: Illustration for Lemma 3.4.

$\sin^{-1}2\beta$. By sine law, $\|p-r\| = \frac{\|p-\tilde{p}\|\cdot\sin\angle p\tilde{p}r}{\sin\angle\tilde{p}rp} \leq \frac{\delta\sin(2\sin^{-1}2\beta)}{\cos(\sin^{-1}2\beta)}$. Note that $\sin(2\sin^{-1}2\beta) \leq 2\sin(\sin^{-1}2\beta) = 4\beta$ and $\cos(\sin^{-1}2\beta) \geq \cos(\sin^{-1}(1/2)) > 0.86$. So $\|p-r\| \leq 4\beta\delta/(0.86) < 5\beta\delta$. By triangle inequality, we get $\|p-q\| \leq \|q-r\|+\|p-r\| = \|\tilde{p}-\tilde{q}\|+\|p-r\| < \|\tilde{p}-\tilde{q}\|+5\beta\delta$. ◻

Lemma 3.5 *Assume that $\beta \leq 1/4$. Let C be any β-cell that lies between the normal segments at the cut-points c_i and c_{i+1}. Then the diameter of C is at most $11\beta f(c_i)$.*

Proof. Let s and t be two points in C. Let p be the projection of s onto a side of C in the direction towards $\tilde{s}$. Similarly, let q be the projection of t onto the same side of C in direction towards $\tilde{t}$. Note that $\tilde{p} = \tilde{s}$ and $\tilde{q} = \tilde{t}$. The triangle inequality and Lemma 3.4 imply that

$$\begin{aligned} \|s-t\| &\leq \|p-q\| + \|p-s\| + \|q-t\| \\ &\leq \|\tilde{p}-\tilde{q}\| + 5\beta\delta + \|p-s\| + \|q-t\|. \end{aligned}$$

Since $\|\tilde{p} - \tilde{q}\| = \|\tilde{s} - \tilde{t}\| \leq 2\beta f(c_i)$ and both $\|p - s\|$ and $\|q - t\|$ are at most $2\beta\delta$, the diameter of C is at most $2\beta f(c_i) + 9\beta\delta \leq 11\beta f(c_i)$. $\square$

3.3 Number of samples in cells and slabs

In this section, we analyze the probabilities of a β-slab and a β-cell containing certain numbers of samples. We first need a lemma that estimates the probability of a sample point lying inside a β-cell and a β-slab.

Lemma 3.6 *Let $\lambda_k = \sqrt{\frac{k^2 \ln^{1+\omega} n}{n}}$ for some positive constant k. Let $r \geq 1$ be a parameter. Let C be a (λ_k/r)-slab or (λ_k/r)-cell. There exist constants κ_1 and κ_2 such that if n is so large that $\lambda_k \leq 1/4$, then $\kappa_2 \lambda_k^2/r^2 \leq \Pr(s \in C) \leq \kappa_1 \lambda_k^2/r^2$.*

Proof. Recall that $L = \int_F \frac{1}{f(x)} dx$. Assume that C lies between the normal segments at the cut-points c_i and c_{i+1}. We use η to denote $F(c_i, c_{i+1})$ as a short hand. By our assumption on λ_k, for any point $x \in \eta$, if C is a λ_k-cell, then $\|x - c_i\| \leq 2\lambda_k f(c_i)/r \leq f(c_i)/2$; if C is a λ_k-slab, then $\|x - c_i\| \leq 2\lambda_k^2 f(c_i)/r^2 \leq f(c_i)/8$. The Lipschitz condition implies that $f(c_i)/2 \leq f(x) \leq 3f(c_i)/2$. If C is a λ_k-slab, then $\Pr(s \in C) = \Pr(\tilde{s}$ lies on $\eta)$, which is $\frac{1}{L} \cdot \int_\eta \frac{1}{f(x)} dx \in [\frac{2\lambda_k^2}{3Lr^2}, \frac{4\lambda_k^2}{Lr^2}]$. If C is λ_k-cell, then $\Pr(\tilde{s}$ lies on $\eta) = \frac{1}{L} \cdot \int_\eta \frac{1}{f(x)} dx \in [\frac{2\lambda_k}{3Lr}, \frac{4\lambda_k}{Lr}]$. Since $\Pr(s \in C \mid \tilde{s}$ lies on $\eta) \in [\lambda_k\delta/(2\delta r), 2\lambda_k\delta/(2\delta r)] = [\lambda_k/(2r), \lambda_k/r]$, $\Pr(s \in C) \in [\frac{\lambda_k^2}{3Lr^2}, \frac{4\lambda_k^2}{Lr^2}]$. $\square$

The following Chernoff bound [8] will be needed.

Lemma 3.7 *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 $\Pr(S_n \geq (1+\sigma)E(S_n)) \leq \exp(-\frac{\sigma^2 E(S_n)}{2(1+\sigma/3)})$.*

We are ready to analyze the probabilities of a β-slab and a β-cell containing certain numbers of samples.

Lemma 3.8 *Let $\lambda_k = \sqrt{\frac{k^2 \ln^{1+\omega} n}{n}}$ for some positive constant k. Let $r \geq 1$ be a parameter. Let C be a (λ_k/r)-slab or (λ_k/r)-cell. Let κ_1 and κ_2 be the constants in Lemma 3.6. Whenever n is so large that $\lambda_k \leq 1/4$, the following hold.*

(i) C is non-empty with probability at least $1 - n^{-\Omega(\ln^\omega n/r^2)}$.

(ii) Assume that $r = 1$. For any constant $\kappa > \kappa_1 k^2$, the number of samples in C is at most $\kappa \ln^{1+\omega} n$ with probability at least $1 - n^{-\Omega(\ln^\omega n)}$.

(iii) Assume that $r = 1$. For any constant $\kappa < \kappa_2 k^2$, the number of samples in C is at least $\kappa \ln^{1+\omega} n$ with probability at least $1 - n^{-\Omega(\ln^\omega n)}$.

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

$$E(S_n) \leq \frac{\kappa_1 n \lambda_k^2}{r^2} = \frac{\kappa_1 k^2 \ln^{1+\omega} n}{r^2}, \qquad E(S_n) \geq \frac{\kappa_2 n \lambda_k^2}{r^2} = \frac{\kappa_2 k^2 \ln^{1+\omega} n}{r^2}.$$

By Lemma 3.7,

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

Consider (ii). Let $\sigma = \frac{\kappa}{\kappa_1 k^2} - 1 > 0$. Since $r = 1$, we have

$$\kappa \ln^{1+\omega} n = \kappa_1 n \lambda_k^2 (1+\sigma) \geq (1+\sigma) E(S_n).$$

By Lemma 3.7,

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

Consider (iii). Let $\sigma = 1 - \frac{\kappa}{\kappa_2 k^2} > 0$. Since $r = 1$, we have

$$\kappa \ln^{1+\omega} n = \kappa_2 n \lambda_k^2 (1-\sigma) \leq (1-\sigma) E(S_n).$$

By Lemma 3.7,

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

□

4 Coarse neighborhood

In this section, we bound the radii of *initial*(s) and *coarse*(s) for each sample s. Then we show that *strip*(s) provides a rough estimate of the slope of the tangent to F at $\tilde{s}$. Recall that $\lambda_k = \sqrt{\frac{k^2 \ln^{1+\omega} n}{n}}$.

4.1 Radius of *initial*(s)

We first need a utility lemma that bounds the distance between two normal segments from below.

Lemma 4.1 *Assume that $\delta \le 1/8$ and $\lambda_k \le 1/4$. Let c_i and c_{i+1} be two consecutive cut-points of a λ_k-partition. For any point on the normal segment at c_{i+1}, its distance from the support line of the normal segment at c_i is at least $|F(c_i, c_{i+1})|/6$.*

Proof. Take any point $q \in F_\alpha$ on the normal segment at c_{i+1}. Let p be the point on F_α such that $\tilde{p} = c_i$. Let r be the orthogonal projection of q onto the tangent at p to F_α. Observe that the distance of q from the support line of the normal segment at c_i is $\|p - r\|$. We are to prove that $\|p - r\| \ge |F(c_i, c_{i+1})|/6$.

For any point $x \in F_\alpha(p, q)$, we use θ_x to denote the non-obtuse angle between the normals at $\tilde{x}$ and c_i. By Lemma 3.3, we have $\theta_x \le 2\sin^{-1}\frac{|F(c_i,c_{i+1})|}{f(c_i)}$. By our assumption on λ_k, $\frac{|F(c_i,c_{i+1})|}{f(c_i)} \le 2\lambda_k^2 < 1/2$. It follows that $\sin^{-1}\frac{|F(c_i,c_{i+1})|}{f(c_i)} < \frac{1.1|F(c_i,c_{i+1})|}{f(c_i)}$. Therefore,

$$\theta_x \le \frac{2.2|F(c_i, c_{i+1})|}{f(c_i)} \tag{1}$$

$$\le 4.4\lambda_k^2. \tag{2}$$

This implies that $F_\alpha(p, q)$ is monotone along the tangent to F_α at p; otherwise, there is a point $x \in F_\alpha(p, q)$ such that $\theta_x = \pi/2 > 4.4\lambda_k^2$, contradiction. It follows that $F(c_i, c_{i+1})$ is also monotone along the tangent to F at c_i.

Refer to Figure 7. Assume that the tangents at p and c_i are horizontal, p lies below c_i, and q lies to the right of p. Let r' be the orthogonal projection of c_{i+1} onto the tangent to F at c_i. Let s (resp. s') be the intersection between the normal at q and the tangent at p (resp. c_i). The monotonicity of $F(c_i, c_{i+1})$ implies that

$$\|c_i - r'\| = \int_{F(c_i,c_{i+1})} \cos\theta_x \, dx \overset{(2)}{\ge} |F(c_i, c_{i+1})| \cdot \cos(4.4\lambda_k^2) > 0.9|F(c_i, c_{i+1})|, \tag{3}$$

as $\cos(4.4\lambda_k^2) \ge \cos(0.275) > 0.9$. Similarly, we get

$$\|p - r\| > 0.9|F_\alpha(p, q)|. \tag{4}$$

If the support line of the normal at c_{i+1} has non-positive slope (see Figure 7(a)), then $\|p - r\| \ge \|c_i - r'\|$. Since $\|c_i - r'\| > 0.9|F(c_i, c_{i+1})|$ by (3), we are done.

Suppose that the support line of the normal at c_{i+1} has positive slope, see Figure 7(b). (Despite the illustration in Figure 7(b), r may not lie between p and s and r' may not lie between c_i and s'.) Starting with triangle inequality, we get

$$\|p - r\| \ge \|p - s\| - \|r - s\| = \|p - s\| - \|q - r\| \cdot \tan\theta_q \tag{5}$$

We are to prove an upper bound for $\|q - r\| \cdot \tan\theta_q$ and a lower bound for $\|p - s\|$. This will yield a lower bound for $\|p - r\|$. By (2), $\theta_q \le 4.4\lambda_k^2 \le 0.275$, so we have

$$\tan\theta_q < 1.03\theta_q \overset{(1)}{<} \frac{3|F(c_i, c_{i+1})|}{f(c_i)}. \tag{6}$$

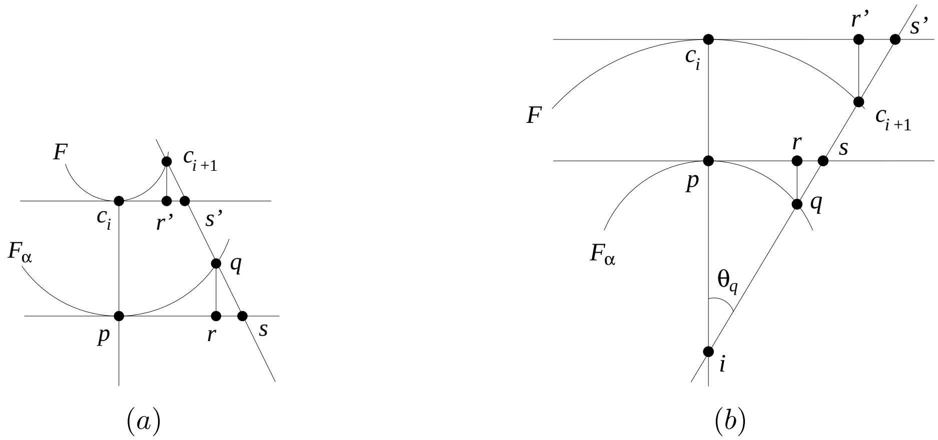

Figure 7: Illustration for Lemma 4.1.

By Lemma 3.4 and our assumption on λ_k, $\|p-q\| \leq \|c_i - c_{i+1}\| + 5\lambda_k\delta \leq 2\lambda_k^2 f(c_i) + 5\lambda_k f(c_i) \leq 1.375 f(c_i)$. Thus, by our assumption that $\delta \leq 1/8$, $\|p-q\| < 2(1-\delta)f(c_i) \leq 2(1-\alpha)f(c_i)$ and so Lemma 3.2(i) applies. We get

$$\|q-r\| \leq \frac{\|p-q\|^2}{2(1-\alpha)f(c_i)} < \frac{\|p-q\|^2}{f(c_i)} \leq \frac{|F_\alpha(p,q)|^2}{f(c_i)}.$$

If $|F_\alpha(p,q)| \geq |F(c_i,c_{i+1})|$, then by (4), $\|p-r\| > 0.9|F(c_i,c_{i+1})|$ and we are done. The remaining case is that $|F_\alpha(p,q)| < |F(c_i,c_{i+1})|$. By our assumption on λ_k, we get

$$\|q-r\| \leq \frac{|F(c_i,c_{i+1})|^2}{f(c_i)} \leq 4\lambda_k^4 f(c_i) < 0.02 f(c_i).$$

Plugging (6) into the above, we obtain

$$\|q-r\| \cdot \tan\theta_q < 0.06|F(c_i,c_{i+1})|. \tag{7}$$

Similarly, we get

$$\begin{aligned} \|c_{i+1} - r'\| &< 0.02 f(c_i), \\ \|r'-s'\| = \|c_{i+1} - r'\| \cdot \tan\theta_q &< 0.06|F(c_i,c_{i+1})|. \end{aligned} \tag{8}$$

Next, we bound $\|p-s\|$ from above. Let i be the intersection point of the normals at c_i and c_{i+1}. Consider the similar triangles ips and ic_is'. We have

$$\|p-s\| = \|c_i - s'\| \cdot \frac{\|p-i\|}{\|c_i - i\|} = \|c_i - s'\| \cdot (1 - \frac{\|p-c_i\|}{\|c_i - i\|}).$$

Observe that $\|p-c_i\| \leq \delta$ and $\|c_i - s'\| = \|c_i - i\| \cdot \tan\theta_q$. Thus,

$$\begin{aligned} \|p-s\| &\geq \|c_i - s'\| \cdot (1 - \frac{\delta\tan\theta_q}{\|c_i - s'\|}) && (9) \\ &= \|c_i - s'\| - \delta\tan\theta_q && (10) \\ &\geq \|c_i - r'\| - \|r'-s'\| - \delta\tan\theta_q && (11) \\ &\overset{(3),(8)\,\&\,(6)}{\geq} 0.84|F(c_i,c_{i+1})| - 3\delta|F(c_i,c_{i+1})|. && (12) \end{aligned}$$

Plugging (12) and (7) into (5), we obtain

$$\begin{aligned}
\|p-r\| &\geq (0.84-3\delta-0.06)|F(c_i,c_{i+1})| \\
&\geq (0.84-0.375-0.06)|F(c_i,c_{i+1})| \\
&> \frac{|F(c_i,c_{i+1})|}{6}.
\end{aligned}$$

□

We are ready to bound the radius of *initial*(s).

Lemma 4.2 *Let h be a constant less than $\sqrt{\frac{1}{3\kappa_1}}$ and let m be a constant greater than $\sqrt{\frac{2}{\kappa_2}}$, where κ_1 and κ_2 are the constants in Lemma 3.6. Let $\psi_h = \lambda_h/3$ and $\psi_m = \sqrt{11\lambda_m}$. Let s be a sample. If $\delta \leq 1/8$, $\lambda_h \leq 1/32$, and $\lambda_m \leq 1/4$, then*

$$\psi_h\sqrt{f(\tilde{s})} \leq \text{radius}(\mathit{initial}(s)) \leq \psi_m\sqrt{f(\tilde{s})}$$

with probability at least $1-O(n^{-\Omega(\ln^\omega n)})$.

Proof. Let D be the disk centered at s that contains $\ln^{1+\omega}$ samples. We first prove the upper bound. Take a λ_m-grid such that s lies on the normal segment at the cut-point c_0. Let C be the λ_m-cell between the normal segments at c_0 and c_1 that contains s. By Lemma 3.8(iii), C contains at least $2\ln^{1+\omega} n$ samples with probability at least $1-n^{-\Omega(\ln^\omega n)}$. Since D contains $\ln^{1+\omega} n$ samples, radius(D) is less than the diameter of C with probability at least $1-n^{-\Omega(\ln^\omega n)}$. By Lemma 3.5, radius$(D) \leq 11\lambda_m f(c_0) = 11\lambda_m f(\tilde{s})$. It follows that radius$(\mathit{initial}(s)) = \sqrt{\text{radius}(D)} \leq \sqrt{11\lambda_m f(\tilde{s})}$.

Next, we prove the lower bound. Take a λ_h-partition such that s lies on the normal segment at the cut-point c_0. Consider the cut-points c_j for $-1 \leq j \leq 1$. (We use c_{-1} to denote the last cut-point picked.) We have $\|c_{-1}-c_0\| \leq |F(c_{-1},c_0)| \leq 2\lambda_h^2 f(c_{-1}) < 0.1 f(c_{-1})$ by our assumption on λ_h. The Lipschitz condition implies that

$$0.9f(c_0) < f(c_{-1}) < 1.2f(c_0). \tag{13}$$

Let ℓ_{-1} and ℓ_1 be the support lines of the normal segments at c_{-1} and c_1. Let d_{-1} and d_1 be the distances from s to ℓ_{-1} and ℓ_1, respectively. We first prove lower bounds on d_{-1} and d_1. By Lemma 4.1,

$$d_{-1} \geq \frac{|F(c_{-1},c_0)|}{6} \geq \frac{\lambda_h^2 f(c_{-1})}{6} \overset{(13)}{>} \frac{\lambda_h^2 f(c_0)}{7}.$$

Assume that $s \in F_\alpha$. Let x be the point on F_α such that $\tilde{x} = c_1$. By Lemma 3.4 and our assumption on λ_h,

$$\begin{aligned}
\|s-x\| &\leq \|c_0-c_1\| + 5\lambda_h\delta \\
&\leq 2\lambda_h^2 f(c_0) + 5\lambda_h f(c_0) \\
&< 0.16f(c_0).
\end{aligned}$$

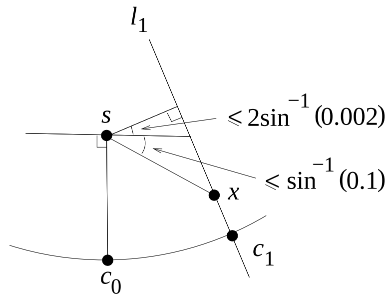

Figure 8: Illustration for Lemma 4.2.

Since $\delta \le 1/8$, $2(1-\alpha) \ge 2(1-\delta) > 0.16$, by Lemma 3.2(ii), $x \in \mathit{cocone}(s, 2\sin^{-1}\frac{0.16}{2(1-\alpha)}) \subseteq \mathit{cocone}(s, 2\sin^{-1}(0.1))$. Since $\|c_0 - c_1\| \le 2\lambda_h^2 f(c_0) < 0.002 f(c_0)$, by Lemma 3.3, the angle between the normal segments at c_0 and c_1 is at most $2\sin^{-1}(0.002)$. Refer to Figure 8. So $d_1 \ge \|s - x\| \cdot \cos(\sin^{-1}(0.1) + 2\sin^{-1}(0.002)) > 0.9 \cdot \|s - x\|$. By Lemma 4.1, $\|s - x\| \ge |F(c_0, c_1)|/6 \ge \lambda_h^2 f(c_0)/6$. We get

$$d_1 > \frac{\lambda_h^2 f(c_0)}{7}.$$

We apply the lower bounds for d_{-1} and d_1 to bound radius($\mathit{initial}(s)$) from below. By Lemma 3.8(ii), the slab between c_{-1} and c_0 and the slab between c_0 and c_1 contain at most $\ln^{1+\omega} n/3$ points each with probability at least $1 - O(n^{-\Omega(\ln^\omega n)})$. Hence, for D to contain $\ln^{1+\omega} n$ points, $\text{radius}(D) > \max\{d_{-1}, d_1\} \ge \lambda_h^2 f(c_0)/7$. Note that $f(\tilde{s}) = f(c_0)$ as $\tilde{s} = c_0$ by construction. It follows that $\text{radius}(\mathit{initial}(s)) = \sqrt{\text{radius}(D)} > \lambda_h \sqrt{f(\tilde{s})}/3$. ◻

4.2 Radius of *coarse*(s)

In this section, we prove an upper bound and a lower bound on the radius of *coarse*(s).

Lemma 4.3 *Assume $\rho \ge 4$ and $\delta \le 1/(25\rho^2)$. Let m be the constant and ψ_m be the parameter in Lemma 4.2. Let s be a sample. If $\lambda_m \le 1/(396\rho^2)$, then*

$$\text{radius}(\mathit{coarse}(s)) \le 5\rho\delta + \psi_m \sqrt{f(\tilde{s})}$$

with probability at least $1 - O(n^{-\Omega(\ln^\omega n)})$.

Proof. Let s_1 and s_2 be points on F_δ^+ and F_δ^- such that $\tilde{s_1} = \tilde{s_2} = \tilde{s}$. Let D be the disk centered at s with radius $5\rho\delta + \psi_m\sqrt{f(\tilde{s})}$. By Lemma 4.2, $\psi_m\sqrt{f(\tilde{s})} \ge \text{radius}(\mathit{initial}(s))$, so D contains $\mathit{initial}(s)$ with probability at least $1 - O(n^{\Omega(\ln^\omega n)})$. We are to show that *coarse*(s) cannot grow beyond D. First, since $\lambda_m \le 1/(396\rho^2)$,

$$\psi_m = \sqrt{11\lambda_m} \le 1/(6\rho) \le 1/24.$$

Observe that both s_1 and s_2 lie inside D. Since $5\rho\delta \le 1/(5\rho) \le 1/20$ and $\psi_m \le 1/24$, $\text{radius}(D) < (1-\delta) f(\tilde{s})$. Thus, the distance between any two points in $D \cap F_\delta^+$ is at most $2(1-\delta)f(\tilde{s})$. By Lemma 3.2(i), the maximum distance between $D \cap F_\delta^+$ and the tangent to F_δ^+ at s_1 is at most $\frac{(5\rho\delta+\psi_m)^2}{2(1-\delta)} < 0.51(5\rho\delta + \psi_m)^2$ as $\delta \le 1/(25\rho^2)$. The same is also true for $D \cap F_\delta^-$. It follows that the samples inside D lie inside a strip of width at most $2\delta+1.1(5\rho\delta+\psi_m)^2 = 2\delta+1.1(5\rho)^2\delta^2+2.2(5\rho)\psi_m\delta+1.1\psi_m^2$. Since $\delta \le 1/(25\rho^2)$ and $\psi_m \le 1/(6\rho)$, we have $1.1(5\rho)^2\delta^2 \le 1.1\delta$, $2.2(5\rho)\psi_m\delta < 1.84\delta$, and $1.1\psi_m^2 < \psi_m/\rho$. We conclude that the strip width is no more than $2\delta + 1.1\delta + 1.84\delta + \psi_m/\rho < 5\delta + \psi_m/\rho \le \text{radius}(D)/\rho$. This shows that *coarse*(s) cannot grow beyond D. ◻

Next, we bound radius(*coarse*(s)) from below. We use $f_{\max}$ to denote $\max_{x\in F} f(x)$.

Lemma 4.4 *Assume that $\delta \le 1/8$ and $\rho \ge 4$. Let h be the constant in Lemma 4.2. Let s be a sample. If $\lambda_h \le 1/32$, then*

$$\text{radius}(\textit{coarse}(s)) \ge \max\{2\sqrt{\rho}\delta, \text{radius}(\textit{initial}(s))\}$$

with probability at least $1 - O(n^{-\Omega(\ln^\omega n/f_{\max})})$.

Proof. Since *coarse*(s) is grown from *initial*(s), radius(*coarse*(s)) $\ge$ radius(*initial*(s)). We are to prove that radius(*coarse*(s)) $\ge 2\sqrt{\rho}\delta$. Let D be the disk that has center s and radius radius(*coarse*(s))$/\sqrt{\rho}$. Let X be the disk centered at $\tilde{s}$ with radius δ. Note that $s \in X$ and X is tangent to F_δ^+ and F_δ^-. Since $\delta \le 1/8$, $f(\tilde{s}) - \delta > \delta$ and so Lemma 3.1 implies that X lies inside the finite region bounded by F_δ^+ and F_δ^-.

Suppose that radius(*coarse*(s)) $< 2\sqrt{\rho}\delta$. Then radius$(D) < 2\delta$. If D contains X, X is a disk inside $D\cap X$ with radius at least radius$(D)/2$. If D does not contain X, then since $s \in X$, $D \cap X$ contains a disk with radius radius$(D)/2$. The width of *strip*(s) is less than or equal to radius(*coarse*(s))$/\rho$ = radius$(D)/\sqrt{\rho}$. Thus, $(D \cap X) -$ *strip*(s) contains a disk Y such that

$$\text{radius}(Y) \ge (\frac{1}{4} - \frac{1}{4\sqrt{\rho}}) \cdot \text{radius}(D) \ge \frac{\text{radius}(D)}{8}.$$

Note that Y is empty and Y lies inside the finite region bounded by F_δ^+ and F_δ^-. Take a point $p \in Y$. Since $p \in Y \subseteq D$ and radius$(D) < 2\delta$, $\|\tilde{p} - \tilde{s}\| \le \|p - \tilde{p}\| + \|s - \tilde{s}\| + \|p - s\| \le 4\delta \le 1/2$ as $\delta \le 1/8$. The Lipschitz condition implies that $f(\tilde{p}) \le 3f(\tilde{s})/2$. Observe that radius$(D)$ = radius(*coarse*(s))$/\sqrt{\rho} \ge$ radius(*initial*(s))$/\sqrt{\rho}$. Thus, Lemma 4.2 implies that radius$(Y) \ge$ radius$(D)/8 \ge \lambda_h\sqrt{f(\tilde{s})}/(24\sqrt{\rho}) > \lambda_h\sqrt{f(\tilde{p})}/(30\sqrt{\rho})$ with probability at least $1 - O(n^{-\Omega(\ln^\omega n)})$. Let $\beta = \lambda_h/(330\sqrt{\rho f_{\max}})$. Then radius$(Y) \ge 11\beta f(\tilde{p})$. By Lemma 3.5, Y contains a β-cell. By Lemma 3.8(i), this β-cell is empty with probability at most $n^{-\Omega(\ln^\omega n/f_{\max})}$. This implies that radius(*coarse*(s)) $< 2\sqrt{\rho}\delta$ occurs with probability at most $O(n^{-\Omega(\ln^\omega n/f_{\max})})$. ◻

4.3 Rough tangent estimate: *strip*(s)

In this section, we prove that the slope of *strip*(s) is a rough estimate of the slope of the tangent at $\tilde{s}$. We first prove a utility lemma about various properties of *coarse*(s) and F_α inside *coarse*(s). Although the lemma contains a long list of properties, their proofs are short.

Lemma 4.5 *Assume $\rho \geq 5$ and $\delta \leq 1/(25\rho^2)$. Let m be the constant and ψ_m be the parameter in Lemma 4.2. Let s be a sample. If $2\sqrt{\rho}\delta \leq \text{radius}(coarse(s)) \leq 5\rho\delta + \psi_m\sqrt{f(\tilde{s})}$ and $\psi_m \leq 1/100$, then for any F_α and for any point $x \in F_\alpha \cap coarse(s)$, the following hold:*

(i) $5\rho\delta + \psi_m \leq 0.05$, $\frac{5\rho\delta+\psi_m}{2(1-\delta)} \leq 0.03$, *and* $\frac{5\rho\delta+\psi_m+2\delta}{2(1-\delta)} \leq 0.03$,

(ii) *$F_\alpha \cap D$ consists of one connected component,*

(iii) *the angle between the normals at s and x is at most* $2\sin^{-1}\frac{5\rho\delta+\psi_m+2\delta}{(1-\delta)} \leq 2\sin^{-1}(0.06)$,

(iv) $x \in cocone(s_1, 2\sin^{-1}\frac{5\rho\delta+\psi_m+2\delta}{2(1-\delta)}) \subseteq cocone(s_1, 2\sin^{-1}(0.03))$ *where s_1 is the point on F_α such that $\tilde{s_1} = \tilde{s}$.*

(v) $0.9f(\tilde{s}) < f(\tilde{x}) < 1.1f(\tilde{s})$,

(vi) *if x lies on the boundary of $coarse(s)$, the distance between s and the orthogonal projection of x onto the tangent at s is at least $0.8 \cdot \text{radius}(coarse(s))$, and*

(vii) *for any $y \in F_\alpha \cap coarse(s)$, the acute angle between xy and the tangent at x is at most* $\sin^{-1}(6\rho\delta + 1.2\psi_m)) \leq \sin^{-1}(0.06)$.

Proof. A straightforward calculation shows (i).

If $F_\alpha \cap D$ consists of more than one connected component, the medial axis of F_α intersects the interior of D. Since F and F_α have the same medial axis, the distance from $\tilde{s}$ to the medial axis is at most $2\,\text{radius}(coarse(s)) \leq 2(5\rho\delta + \psi_m\sqrt{f(\tilde{s})}) \leq 2(5\rho\delta + \psi_m)f(\tilde{s}) < f(\tilde{s})$ by (i), contradiction. This proves (ii).

Let s_1 be the point on F_α such that $\tilde{s_1} = \tilde{s}$. The distance $\|s_1 - x\| \leq \|s - x\| + \|s - s_1\| \leq (5\rho\delta + \psi_m + 2\delta)f(\tilde{s})$. By Lemma 3.3, the angle between the normals at s_1 and x is at most $2\sin^{-1}\frac{\|s_1-x\|}{(1-\delta)f(\tilde{s})} \leq 2\sin^{-1}\frac{5\rho\delta+\psi_m+2\delta}{(1-\delta)} \leq 2\sin^{-1}(0.06)$ by (i). This proves (iii).

By Lemma 3.2(ii), $x \in cocone(s_1, 2\sin^{-1}\frac{\|s_1-x\|}{2(1-\delta)f(\tilde{s})}) \subseteq cocone(s_1, 2\sin^{-1}(0.03))$. This proves (iv).

The distance $\|\tilde{s} - \tilde{x}\| \leq \|s - \tilde{s}\| + \|s - x\| + \|x - \tilde{x}\| \leq (5\rho\delta + \psi_m + 4\delta)f(\tilde{s}) < 0.1f(\tilde{s})$. Then the Lipschitz condition implies (v).

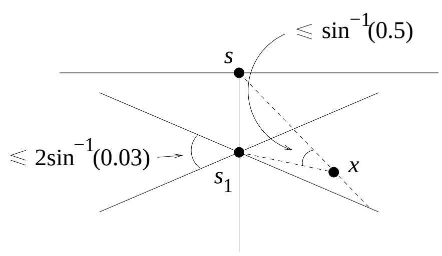

Figure 9:

Consider (vi). Refer to Figure 9. Assume that the tangent at s is horizontal. By sine law, $\sin\angle sxs_1 = \frac{\|s-s_1\|\cdot\sin\angle ss_1x}{\|s-x\|} \leq \frac{2\delta}{\text{radius}(coarse(s))}$ as $\|s - s_1\| \leq 2\delta$ and $\|s - x\| = \text{radius}(coarse(s))$. Since $\text{radius}(coarse(s)) \geq 2\sqrt{\rho}\delta$ and $\rho \geq 4$, we have $\angle sxs_1 \leq \sin^{-1}\frac{1}{\sqrt{\rho}} \leq \sin^{-1}(0.5)$. By (iv),

$\angle s_1 s x \geq \pi - \angle s x s_1 - (\pi/2 + \sin^{-1}(0.03)) \geq \pi/2 - \sin^{-1}(0.5) - \sin^{-1}(0.03)$. Thus, the horizontal distance between s and x is equal to $\|s-x\| \cdot \sin \angle s_1 s x \geq \|s-x\| \cdot \cos(\sin^{-1}(0.5) + \sin^{-1}(0.03)) > 0.8 \cdot \|s - x\|$.

Consider (vii). Since $y \in F_\alpha \cap \mathit{coarse}(s)$, $\|x-y\| \leq 2\,\text{radius}(\mathit{coarse}(s)) \leq 2(5\rho\delta + \psi_m) f(\tilde{s}) < 0.1 f(\tilde{s})$ by (i). So Lemma 3.2(ii) applies and the acute angle between xy and the tangent at x is at most $\sin^{-1} \frac{\|x-y\|}{2(1-\delta)f(\tilde{x})} \leq \sin^{-1} \frac{(5\rho\delta+\psi_m)f(\tilde{s})}{(1-\delta)f(\tilde{x})}$. Since $f(\tilde{x}) \geq 0.9 f(\tilde{s})$ by (v) and $\delta \leq 1/(25\rho^2)$, the acute angle is less than $\sin^{-1}(1.2(5\rho\delta + \psi_m))$, which is less than $\sin^{-1}(0.06)$ by (i). □

We are ready to analyze the slope of *strip*(s). We highlight the key ideas before giving the proof. Let $\mathcal{B}$ be the region between F_δ^+ and F_δ^- inside *coarse*(s). If *strip*(s) makes a large angle with the tangent at $\tilde{s}$, *strip*(s) would cut through $\mathcal{B}$ in the middle. In this case, if $\mathcal{B} \cap \mathit{strip}(s)$ is narrow, there would be a lot of areas in $\mathcal{B}$ outside *strip*(s). But these areas must be empty which occur with low probability. Otherwise, if $\mathcal{B} \cap \mathit{strip}(s)$ is wide, we show that *strip*(s) can be rotated to reduce its width further, contradiction. We give the detailed proof below.

Lemma 4.6 *Assume that $\rho \geq 5$ and $\delta \leq 1/(25\rho^2)$. Let m be the constant and ψ_m be the parameter in Lemma 4.2. Let s be a sample. For sufficiently large n, the acute angle between the tangent at $\tilde{s}$ and the direction of strip(s) is at most* $3\sin^{-1} \frac{5\rho\delta+\psi_m+2\delta}{(1-\delta)} + \sin^{-1}(6\rho\delta + 1.2\psi_m) \leq 4\sin^{-1}(0.06)$ *with probability at least* $1 - O(n^{-\Omega(\ln^\omega n/f_{\max})})$.

Proof. Let ℓ_1 and ℓ_2 be the lower and upper bounding lines of *strip*(s). Without loss of generality, we assume that the normal at $\tilde{s}$ is vertical, the slope of *strip*(s) is non-negative, $F_\delta^- \cap \mathit{coarse}(s)$ lies below $F_\delta^+ \cap \mathit{coarse}(s)$, and $\psi_m \leq 1/100$ for sufficiently large n. Let h and m be the constants and ψ_h and ψ_m be the parameters in Lemma 4.2. We first assume that $\max\{2\sqrt{\rho}\delta, \psi_h\sqrt{f(\tilde{s})}\} \leq \text{radius}(\mathit{coarse}(s)) \leq 5\rho\delta + \psi_m\sqrt{f(\tilde{s})}$ and take the probability of its occurrence into consideration later. As a short hand, we use η_1 to denote $\frac{5\rho\delta+\psi_m+2\delta}{(1-\delta)}$ and η_2 to denote $6\rho\delta + 1.2\psi_m$.

Observe that both ℓ_1 and ℓ_2 must intersect the space that lies between F_δ^+ and F_δ^- inside *coarse*(s). Otherwise, we can squeeze *strip*(s) and reduce its width, contradiction. If ℓ_1 intersects $F_\alpha \cap \mathit{coarse}(s)$ twice for some α, then ℓ_1 is parallel to the tangent at some point on $F_\alpha \cap \mathit{coarse}(s)$. By Lemma 4.5(iii), the direction of *strip*(s) makes an angle at most $2\sin^{-1}\eta_1$ with the horizontal and we are done. Similarly, we are done if ℓ_2 intersects $F_\alpha \cap \mathit{coarse}(s)$ twice for some α. The remaining case is that both ℓ_1 and ℓ_2 intersect $F_\alpha \cap \mathit{coarse}(s)$ for any α at most once. Suppose that the acute angle between the direction of *strip*(s) and the horizontal is more than $3\sin^{-1}\eta_1 + \sin^{-1}\eta_2$. We show that this occurs with probability $O(n^{-\Omega(\ln^\omega n)})$.

Let q be the right intersection point between F_δ^- and the boundary of *coarse*(s). If ℓ_1 intersects $F_\delta^- \cap \mathit{coarse}(s)$, let p denote the intersection point; otherwise, let p denote the leftmost intersection point between F_δ^- and the boundary of *coarse*(s). Refer to Figure 10(a). We claim that $F_\delta^-(p,q)$ lies below ℓ_1. If ℓ_1 does not intersect $F_\delta^- \cap \mathit{coarse}(s)$, then this is clearly true. Otherwise, by Lemma 4.5(iii), the magnitude of the slope of the tangent at p is at most $2\sin^{-1}\eta_1$. Since the slope of ℓ_1 is more than $3\sin^{-1}\eta_1 + \sin^{-1}\eta_2$, F_δ^- crosses ℓ_1 at p from above to below. So $F_\delta^-(p,q)$ lies below ℓ_1.

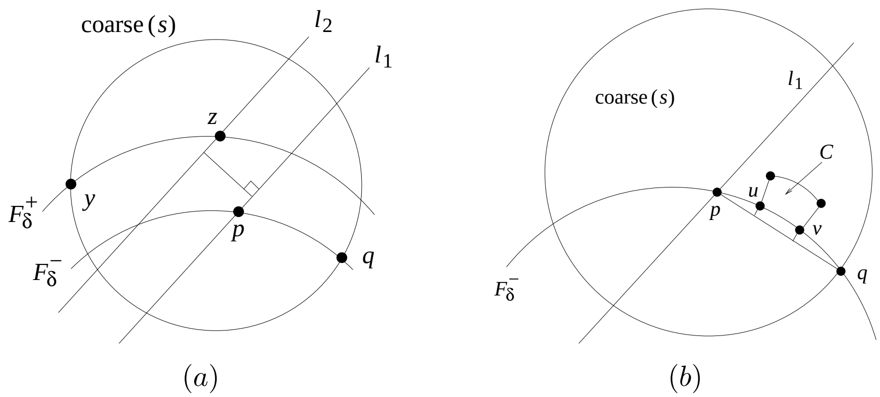

Figure 10:

We show that $\|p-q\| \leq \psi_h\sqrt{f(\tilde{s})}/2$ with probability at least $1-n^{-\Omega(\ln^\omega n)}$. Notice that pq is parallel to the tangent to F^-_δ at some point on $F^-_\delta(p,q)$. By Lemma 4.5(iii), the tangent to $F^-_\delta(p,q)$ turns by an angle at most $4\sin^{-1}(0.06) < \pi/2$ from p to q. This implies that $F^-_\delta(p,q)$ is monotone with respect to the perpendicular direction of pq.

We divide pq into three equal segments. Let u and v be the intersection points between $F^-_\delta(p,q)$ and the perpendiculars of pq at the dividing points. Assume that v follows u along $F^-_\delta(p,q)$. Refer to Figure 10(b). Suppose that $\|p-q\| > \psi_h\sqrt{f(\tilde{s})}/2$. Then

$$|F^-_\delta(u,v)| \geq \frac{\|p-q\|}{3} \geq \frac{\psi_h\sqrt{f(\tilde{s})}}{6}. \tag{14}$$

Since $f(\tilde{u}) < 1.1f(\tilde{s})$ by Lemma 4.5(v), $|F^-_\delta(u,v)| > \psi_h f(\tilde{u})/7$. Consider a λ_k-grid where $k = h/231$ and $\tilde{u}$ is a cut-point. (Note that $\lambda_k = \psi_h/77$.) Let C be the λ_k-cell that touches $F^-_\delta(u,v)$ and the normal segment through u. By Lemma 3.5, the diameter of C is at most $11\lambda_k f(\tilde{u}) = \psi_h f(\tilde{u})/7 < |F^-_\delta(u,v)|$. So the bottom side of C lies inside $F^-_\delta(u,v)$. Let $\mathcal{R}$ be the region inside *coarse*(s) that lies below ℓ_1 and above $F^-_\delta(p,q)$. From any point $x \in F^-_\delta(u,v)$, if we shoot a ray along the normal at x into $\mathcal{R}$, either the ray will leave C first or the ray will hit ℓ_1 or the boundary of *coarse*(s) in $\mathcal{R}$. We are to prove that the distances from x to ℓ_1 and the boundary of *coarse*(s) in $\mathcal{R}$ are more than $2\lambda_k\delta$. This shows that the ray always leaves C first, so C lies completely inside *coarse*(s) and below ℓ_1. Then the upper bound on $\|p-q\|$ follows as C is empty with probability at most $n^{-\Omega(\ln^\omega n)}$ by Lemma 3.8(i).

Consider the distance from x to ℓ_1. By Lemma 4.5(iii), the angle between ℓ_1 and the tangent at p (measured by rotating ℓ_1 in the clockwise direction) is at least $3\sin^{-1}\eta_1 + \sin^{-1}\eta_2 - 2\sin^{-1}\eta_1 = \sin^{-1}\eta_1 + \sin^{-1}\eta_2$ and at most $\pi/2 + 2\sin^{-1}\eta_1$. By Lemma 4.5(vii), the acute angle between px and the tangent at p is at most $\sin^{-1}\eta_2$. So the angle between px and ℓ_1 is at least $\sin^{-1}\eta_1$ and at most $\pi/2 + 2\sin^{-1}\eta_1 + \sin^{-1}\eta_2$. This implies that the distance from x to ℓ_1 is at least $\|p-x\| \cdot \min\{\eta_1\,,\, \cos(2\sin^{-1}\eta_1 + \sin^{-1}\eta_2)\}$. By Lemma 4.5(i), $\eta_1 \leq 0.06 < \cos(3\sin^{-1}(0.06)) \leq \cos(2\sin^{-1}\eta_1 + \sin^{-1}\eta_2)$. Therefore, the distance from x to ℓ_1 is at least $\|p-x\| \cdot \eta_1 > 5\rho\delta \cdot \|p-x\| \geq 20\delta \cdot (\|p-q\|/3) \overset{(14)}{>} 3\delta\psi_h\sqrt{f(\tilde{s})}$. Since $\lambda_k = \psi_h/88$, this distance is greater than $2\lambda_k\delta$.

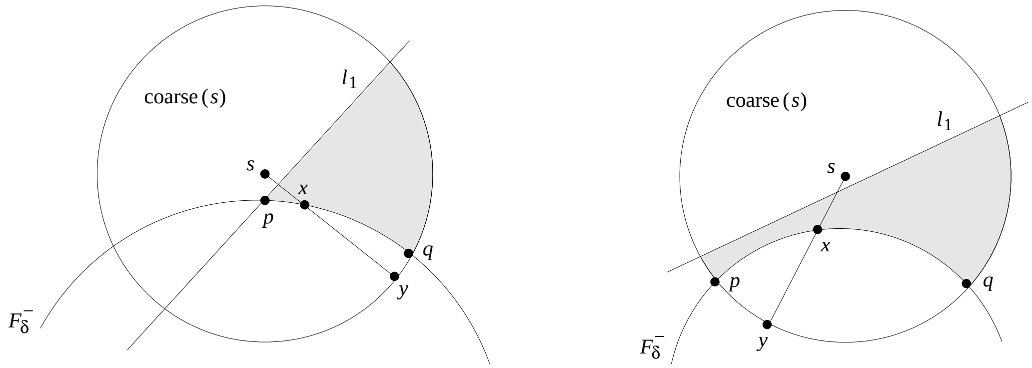

(a) (b)

Figure 11:

Next, we consider the distance d from x to the boundary of $coarse(s)$ in $\mathcal{R}$. Take a radius sy of $coarse(s)$ that passes through x. Suppose that sy intersects $F_\delta^- \cap coarse(s)$ only once at x. Refer to Figure 11. In this case, xy lies outside $\mathcal{R}$. Therefore, if ℓ_1 intersects $F_\delta^- \cap coarse(s)$ at p (Figure 11(a)), then $d = \|q - x\|$; if ℓ_1 does not intersect $F_\delta^- \cap coarse(s)$ (Figure 11(b)), then $d = \min\{\|p - x\|, \|q - x\|\}$. By (14), $d \geq \|p - q\|/3 \geq \psi_h \sqrt{f(\tilde{s})}/6 > 2\lambda_k \delta$. The remaining possibility is that sy intersects $F_\delta^- \cap coarse(s)$ more than once. Then xy is parallel to the tangent at some point on $F_\delta^- \cap coarse(s)$. By Lemma 4.5(iii), the acute angle between xy and the tangent at x is at most $4\sin^{-1}\eta_1$. By Lemma 4.5(vii), the acute angle between qx and the tangent at x is at most $\sin^{-1}\eta_2$. So the angle between qx and xy is at most $4\sin^{-1}\eta_1 + \sin^{-1}\eta_2$. It follows that $d \geq \|x - y\| \geq \|q - x\| \cdot \cos(4\sin^{-1}\eta_1 + \sin^{-1}\eta_2) \geq \|q - x\| \cdot \cos(5\sin^{-1}(0.08)) > 0.9 \cdot \|q - x\| \geq 0.9 \cdot (\|p - q\|/3) \geq 0.15\psi_h > \sqrt{f(\tilde{s})} > 2\lambda_k\delta$.

In all, C lies below ℓ_1 and inside $coarse(s)$. So C must be empty which occurs with probability at most $n^{-\Omega(\ln^\omega n)}$ by Lemma 3.8(i). It follows that $\|p - q\| \leq \psi_h \sqrt{f(\tilde{s})}/2$ with probability at least $1 - n^{-\Omega(\ln^\omega n)}$. By Lemma 4.5(vi), the horizontal distance between q and the left intersection point between F_δ^- and the boundary of $coarse(s)$ is at least $1.6 \cdot \text{radius}(coarse(s)) \geq 1.6\psi_h \sqrt{f(\tilde{s})} > \|p - q\|$. We conclude that p lies on $F_\delta^- \cap coarse(s)$, which implies that ℓ_1 intersects $F_\delta^- \cap coarse(s)$ exactly once at p.

Refer to Figure 10(a) and Figure 12. Let y be the leftmost intersection point between F_δ^+ and the boundary of $coarse(s)$. Symmetrically, we can also show that ℓ_2 intersects $F_\delta^+ \cap coarse(s)$ exactly once at some point z, $F_\delta^+(y, z)$ lies above ℓ_2, and $\|y - z\| \leq \psi_h \sqrt{f(\tilde{s})}/2$ with probability at least $1 - n^{-\Omega(\ln^\omega n)}$.

Consider the projections of $F_\delta^+(y, z)$ and $F_\delta^-(p, q)$ onto the horizontal diameter of $coarse(s)$ through s. By Lemma 4.5(vi), the projections of y and q are at distance at least $0.8 \cdot \text{radius}(coarse(s))$ from s. Thus, the distance between the projections of $F_\delta^+(y, z)$ and $F_\delta^-(p, q)$ is at least $1.6 \cdot \text{radius}(coarse(s)) - \|p - q\| - \|y - z\| \geq 1.6 \cdot \text{radius}(coarse(s)) - \psi_h \sqrt{f(\tilde{s})} \geq 1.6 \cdot \text{radius}(coarse(s)) - \text{radius}(coarse(s)) > \text{radius}(coarse(s))/\rho$. That is, this distance is greater than the width of $strip(s)$. But then we can rotate ℓ_1 and ℓ_2 around p and z, respectively, in the clockwise direction to reduce the width of $strip(s)$ while not losing any

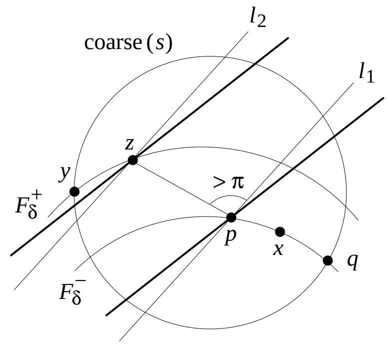

Figure 12:

sample inside *coarse*(s). See Figure 12. This is impossible. This implies that, under the condition that $\max\{2\sqrt{\rho}\delta, \psi_h\sqrt{f(\tilde{s})}\} \leq \text{radius}(coarse(s)) \leq 5\rho\delta + \psi_m\sqrt{f(\tilde{s})}$, the acute angle between the direction of *strip*(s) and the tangent at $\tilde{s}$ is at most $3\sin^{-1}\eta_1 + \sin^{-1}\eta_2$ with probability at least $1 - O(n^{\Omega(\ln^\omega n)})$. By Lemmas 4.2, 4.3, and 4.4, the inequalities $\max\{2\sqrt{\rho}\delta, \psi_h\sqrt{f(\tilde{s})}\} \leq \text{radius}(coarse(s)) \leq 5\rho\delta + \psi_m\sqrt{f(\tilde{s})}$ hold with probability at least $1 - O(n^{\Omega(\ln^\omega n)/f_{\max}})$. So the lemma follows. □

5 Guarantees

In this section, we prove that the reconstruction returned by our algorithm is faithful with high probability. We first prove the pointwise convergence. Then we prove that the reconstruction is homeomorphic to the true curve. Afterwards, we combine these results to prove our main result in this paper.

5.1 Pointwise convergence

Recall that our algorithm computes a center point for each sample. Eventually, a subset of these center points become the vertices of the output curve. Our goal is to show that all center points converge to F as n tends to ∞. To this end, we show that our algorithm aligns *refined*(s) approximately well with the normal at $\tilde{s}$. Then we prove the pointwise convergence. (See Lemmas 5.3 and 5.4.) We first prove two utility lemmas, Lemmas 5.1 and 5.2.

5.1.1 Utility lemmas

Recall that we rotate *refined*(s) in the clockwise and anti-clockwise directions to estimate the normal at $\tilde{s}$. The range of rotation is $[0, \pi/10]$. Let θ_s be the angle between the upward direction of *refined*(s) and the upward normal at $\tilde{s}$. If the upward direction of *refined*(s) points to the left of the upward normal at $\tilde{s}$, θ_s is positive. Otherwise, θ_s is negative. For any F_α and for any point $p \in F_\alpha \cap$ *refined*(s), let γ_p be the angle between the upward direction of *refined*(s) and the upward normal at $\tilde{p}$. The sign of γ_p is determined in the same way as θ_s.

Lemma 5.1 *Assume that $\delta \leq 1/(25\rho^2)$ and $\rho \geq 5$. Let s be a sample. Assume that refined(s) is rotated within an angle of $\pi/10$. Let $W_s = \mathrm{width}(\mathit{refined}(s))$. For sufficiently large n, the following hold throughout the rotation with probability at least $1 - O(n^{-\Omega(\ln^\omega n/f_{\max})})$.*

(i) $W_s \leq 0.1 f(\tilde{s})$.

(ii) $\theta_s \in [-\pi/5, \pi/5]$ *and* $\theta_s = 0$ *at some point during the rotation.*

(iii) Any line, which is parallel to refined(s) and inside refined(s), intersects $F_\alpha \cap$ coarse(s) for any α exactly once.

(iv) For any F_α and for any point $p \in F_\alpha \cap$ refined(s), $\theta_s - 0.2|\theta_s| - 3W_s/f(\tilde{s}) \leq \gamma_b \leq \theta_s + 0.2|\theta_s| + 3W_s/f(\tilde{s})$.

Proof. We first assume that $\max\{2\sqrt{\rho}\delta, \psi_h\sqrt{f(\tilde{s})}\} \leq \mathrm{radius}(\mathit{coarse}(s)) \leq 5\rho\delta + \psi_m\sqrt{f(\tilde{s})}$ and $\mathrm{radius}(\mathit{initial}(s)) \leq \psi_m\sqrt{f(\tilde{s})}$. We will take the consideration of the probabilities of their occurrences later.

Since $W_s \leq \sqrt{\mathrm{radius}(\mathit{initial}(s))} \leq \sqrt{\psi_m} f(\tilde{s})^{1/4}$ and $\psi_m \leq 0.01$ for sufficiently large n, $W_s \leq 0.1 f(\tilde{s})$. This proves (i).

By Lemma 4.6, for sufficiently large n, the acute angle between the normal at $\tilde{s}$ and the initial *refined*(s) is at most $4\sin^{-1}(0.06) < \pi/10$. Since the range of rotation is $[0, \pi/10]$, $\theta_s \in [-\pi/5, \pi/5]$ and $\theta_s = 0$ at some point during the rotation. This proves (ii).

Consider (iii). Let ℓ be a line that is parallel to *refined*(s) and inside *refined*(s). We first prove that ℓ intersects F_α. Refer to Figure 13. Without loss of generality, assume that the normal at $\tilde{s}$ is vertical, the slope of *refined*(s) is positive, and ℓ is below s. Let s_1 and s_2 be the points on F_δ^+ and F_δ^-, respectively, such that $\tilde{s_1} = \tilde{s_2} = \tilde{s}$. Shoot two rays upward from s_1 with slopes $\pm\sin^{-1}(0.03)$. Also, shoot two rays downward from s_2 with slopes $\pm\sin^{-1}(0.03)$. Let $\mathcal{R}$ be the region inside *coarse*(s) bounded by these four rays. By Lemma 4.5(iv), $F_\alpha \cap$ *coarse*(s) lies inside $\mathcal{R}$. Let x be the upper right vertex of $\mathcal{R}$. Let y be the right endpoint of a horizontal chord through s_1. Let L be the line that passes through x and is parallel to ℓ. Let L' be the line that passes through s and is parallel to ℓ. Let z be the point on L such that $s_1 z$ is perpendicular to L.

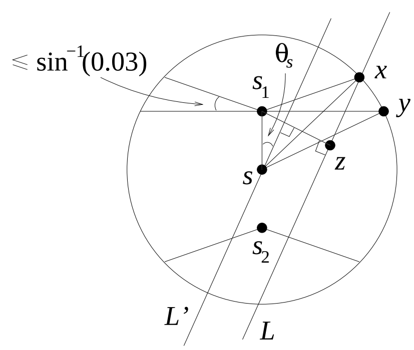

Figure 13:

We claim that L' is above L and L and L' intersect both the upper and lower boundaries of $\mathcal{R}$. By Lemma 4.5(iv), $\angle xs_1y \le \sin^{-1}(0.03)$, so $\angle xsy \le 2\sin^{-1}(0.03)$. Observe that $\cos\angle s_1sy = \frac{\|s-s_1\|}{\|s-y\|} \le \frac{2\delta}{\text{radius}(\mathit{coarse}(s))}$. Since radius$(\mathit{coarse}(s)) \ge 2\sqrt{\rho}\delta$, $\cos\angle s_1sy \le 1/\sqrt{\rho} \le 1/\sqrt{5}$ which implies that $\angle s_1sy > \pi/3$. Since $\angle s_1sx = \angle s_1sy - \angle xsy$, we get

$$\angle s_1sx \ge \pi/3 - 2\sin^{-1}(0.03) > \pi/5 \ge \theta_s. \tag{15}$$

So L' cuts through the angle between ss_1 and sx. It follows that L' is above L. Observe that L' intersects s_1x. By symmetry, L' intersects the left downward ray from s_2 too. We conclude that L and L' intersect both the upper and lower boundaries of $\mathcal{R}$.

Since $\theta_s \le \pi/5$ and $\angle sxz = \angle s_1sx - \theta_s$, by (15), $\angle sxz \ge \pi/3 - 2\sin^{-1}(0.03) - \pi/5 > 0.3$. The distance from s to L is equal to $\|s-x\|\cdot\sin\angle sxz > \|s-x\|\cdot\sin(0.3) > 0.2\cdot\text{radius}(\mathit{coarse}(s))$. Recall that ℓ lies below s by our assumption. The distance between ℓ and s is at most $W_s/2$ and our algorithm enforces that $W_s/2 \le \text{radius}(\mathit{coarse}(s))/6$. So ℓ lies between L' and L. Since L and L' intersect both the upper and lower boundaries of $\mathcal{R}$, so does ℓ. It follows that ℓ must intersect $F_\alpha \cap \mathit{coarse}(s)$.

Next, we show that ℓ intersects $F_\alpha \cap \mathit{coarse}(s)$ exactly once. If not, ℓ is parallel to the tangent at some point on $F_\alpha \cap \mathit{coarse}(s)$. By Lemma 4.5(iii), the angle between ℓ and the vertical is at least $\pi/2 - 2\sin^{-1}(0.06) > \pi/5$, contradicting the fact that $|\theta_s| \le \pi/5$.

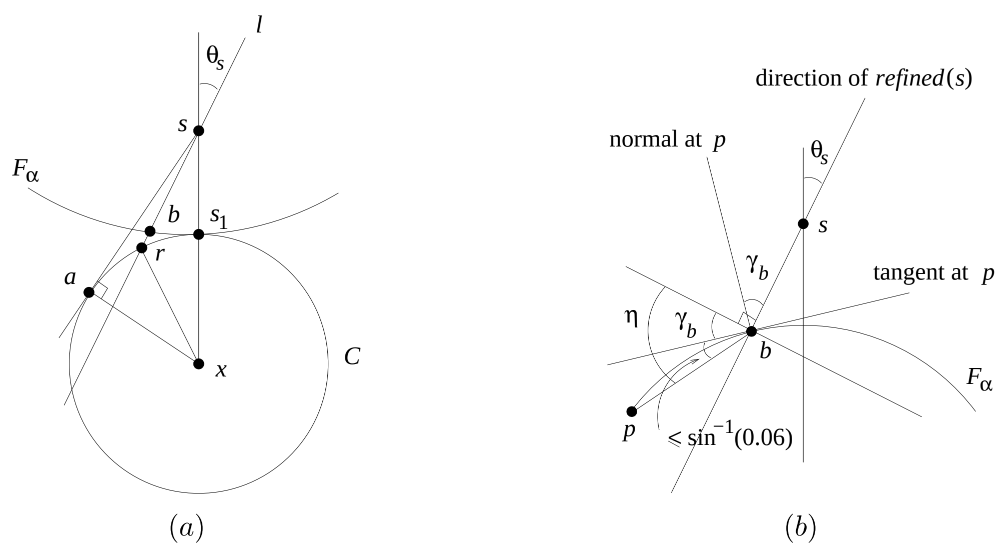

Figure 14:

Consider (iv). Let ℓ be a line that is parallel to *refined*(s) and passes through s. By (iii), ℓ intersects F_α at some point b. We first prove that $\theta_s - 0.2|\theta_s| \le \gamma_b \le \theta_s + 0.2|\theta_s|$. Let s_1 be the point on F_α such that $\tilde{s} = \tilde{s_1}$. Assume that the tangent at s is horizontal, s is above s_1, and b is to the left of s. Let C be the circle tangent to F_α at s_1 that lies below s_1, is centered at x, and has radius $f(\tilde{s}) - \delta$. By Lemma 3.1, F_α does not intersect the interior of C. Refer to Figure 14(a). Let sa be a tangent to C that lies on the left of x. We claim that $\angle asx > |\theta_s|$. Otherwise, $\|s-x\| \ge \|a-x\|/\sin(\pi/5) = (f(\tilde{s})-\delta)/\sin(\pi/5) > f(\tilde{s})+\delta \ge \|s-x\|$, contradiction. So sb lies between sa and sx. Let sr be the extension of sb such that r lies on C. We have $\|a-s\| = \sqrt{\|s-x\|^2 - \|a-x\|^2} \le \sqrt{(f(\tilde{s})+\delta)^2 - (f(\tilde{s})-\delta)^2} = 2\sqrt{\delta f(\tilde{s})}$. Thus, $\|r-s\| \le \|a-s\| \le 2\sqrt{\delta f(\tilde{s})}$. Observe that

$$\angle rxs = \sin^{-1}\frac{\|r-s\|\cdot\sin|\theta_s|}{\|r-x\|} \le \sin^{-1}\frac{2\sqrt{\delta f(\tilde{s})}\cdot|\theta_s|}{\|r-x\|}.$$

Since $\delta \le 1/(25\rho^2)$ and $|\theta_s| \le \pi/5$, we have

$$\frac{2\sqrt{\delta f(\tilde{s})}\cdot|\theta_s|}{\|r-x\|} = \frac{2\sqrt{\delta f(\tilde{s})}\cdot|\theta_s|}{f(\tilde{s})-\delta} \le \frac{2\sqrt{\delta}\cdot|\theta_s|}{1-\delta} < 0.06. \tag{16}$$

Combing (16) with the following fact

$$x \le 0.6 \;\Rightarrow\; \sin^{-1} x < 1.1x, \tag{17}$$

we get $\angle rxs < \frac{2.2\sqrt{\delta f(\tilde{s})}\cdot|\theta_s|}{\|r-x\|}$. Since $\|b-s_1\| \le \|r-s_1\| = \|r-x\|\cdot 2\sin\frac{\angle rxs}{2}$, we get

$$\|b-s_1\| \le \|r-x\|\cdot\angle rxs \le 2.2\sqrt{\delta f(\tilde{s})}\cdot|\theta_s|.$$

Let γ' be the acute angle between the normals at b and s_1. By Lemma 3.3, $\gamma' \le 2\sin^{-1}\frac{\|b-s_1\|}{(1-\alpha)f(\tilde{s})} \le 2\sin^{-1}\frac{2.2\sqrt{\delta}\cdot|\theta_s|}{1-\alpha} \le 2\sin^{-1}\frac{2.2\sqrt{\delta}\cdot|\theta_s|}{1-\delta}$. By (16) and (17), we conclude that $\gamma' < \frac{4.84\sqrt{\delta}\cdot|\theta_s|}{1-\delta} < 0.2|\theta_s|$. It follows that

$$\theta_s - 0.2|\theta_s| \le \theta_s - \gamma' \le \gamma_b \le \theta_s + \gamma' \le \theta_s + 0.2|\theta_s|.$$

Next, we prove the upper and lower bounds for γ_p for any point $p \in F_\alpha \cap \mathit{refined}(s)$. Let η be the acute angle between bp and the line that passes through b and is perpendicular to $\mathit{refined}(s)$. See Figure 14(b). By Lemma 4.5(vii), the acute angle between bp and the tangent at b is at most $\sin^{-1}(0.06)$. It follows that $\eta \le \gamma_b + \sin^{-1}(0.06) \le \theta_s + 0.2|\theta_s| + \sin^{-1}(0.06) \le 1.2(\pi/5) + \sin^{-1}(0.06) < 0.9$. Thus,

$$\|b-p\| \le \frac{W_s}{2\cos\eta} < 0.9W_s.$$

Note that $W_s \le \text{radius}(\mathit{coarse}(s))/3 \le (5\rho\delta + \psi_m)f(\tilde{s})/3$, which is less than $0.02f(\tilde{s})$ by Lemma 4.5(i). Also, by Lemma 4.5(v), $f(\tilde{p}) \ge 0.9f(\tilde{s})$. It follows that

$$\|b-p\| < 0.9W_s \le 0.02f(\tilde{p}). \tag{18}$$

So we can invoke Lemma 3.3 to bound the angle γ'' between the normals at b and p:

$$\gamma'' \le 2\sin^{-1}\frac{\|b-p\|}{(1-\alpha)f(\tilde{p})} \le 2\sin^{-1}\frac{0.9W_s}{(1-\alpha)f(\tilde{p})} \le 2\sin^{-1}\frac{W_s}{f(\tilde{p})}.$$

By (18), $W_s/f(\tilde{p}) < 0.03$. So by (17), we get $\gamma'' \le 2.2W_s/f(\tilde{p})$. Since $f(\tilde{p}) \ge 0.9f(\tilde{s})$, we conclude that $\gamma'' < 3W_s/f(\tilde{s})$. This implies that

$$\theta_s - 0.2|\theta_s| - 3W_s/f(\tilde{s}) \le \gamma_b - \gamma'' \le \gamma_p \le \gamma_b + \gamma'' \le \theta_s + 0.2|\theta_s| + 3W_s/f(\tilde{s}).$$

Finally, we have proved the lemma under the conditions that $\max\{2\sqrt{\rho}\delta, \psi_h\sqrt{f(\tilde{s})}\} \leq \text{radius}(\textit{coarse}(s)) \leq 5\rho\delta + \psi_m\sqrt{f(\tilde{s})}$ and $\text{radius}(\textit{initial}(s)) \leq \psi_m\sqrt{f(\tilde{s})}$. These conditions hold with probabilities at least $1 - O(n^{-\Omega(\ln^\omega n/f_{\max})})$ by Lemmas 4.2, 4.3, and 4.4. So the lemma follows. ◻

Lemma 5.2 *Assume that $\delta \leq 1/(25\rho^2)$ and $\rho \geq 5$. Let s be a sample. Let H be a strip that is parallel to refined(s) and lies inside refined(s). For any F_α and for any two points u and v on on $F_\alpha \cap H$, whenever n is sufficiently large, the following hold with probability at least $1 - O(n^{-\Omega(\ln^\omega n/f_{\max})})$.*

(i) $\|u - v\| < 3\,\text{width}(H)$.

(ii) The angle between the normals at u and v is at most $9\,\text{width}(H)$.

(iii) The acute angle between uv and the tangent to F_α at u is at most $5\,\text{width}(H)$.

Proof. Let ϕ be the acute angle between uv and the tangent to F_α at u. Let η be the acute angle between uv and the direction of *refined*(s). By Lemma 4.5(vii), $\phi \leq \sin^{-1}(0.06)$. So $\eta \geq \pi/2 - \gamma_u - \phi \geq \pi/2 - \gamma_u - \sin^{-1}(0.06)$. By Lemma 5.1(i), (ii), and (iv), $\eta \geq \pi/2 - 1.2(\pi/5) - 3(0.1) - \sin^{-1}(0.06) > 0.4$. Thus, $\|u - v\| \leq \frac{\text{width}(H)}{\sin\eta} \leq \frac{\text{width}(H)}{\sin(0.4)} < 3\,\text{width}(H)$. This proves (i).

Consider (ii). Note that $W_s \leq \text{radius}(\textit{coarse}(s))/3 \leq (5\rho\delta + \psi_m)f(\tilde{s})/3$. So by (i), $\|u - v\| \leq 3W_s \leq (5\rho\delta + \psi_m)f(\tilde{s})$. By Lemma 4.5(i) and (v), $5\rho\delta + \psi_m \leq 0.05$ and $f(\tilde{u}) \geq 0.9f(\tilde{s})$. It follows that

$$\|u - v\| < 0.06f(\tilde{u}). \tag{19}$$

Thus, we can invoke Lemma 3.3 to bound the angle ξ between the normals at u and v:

$$\xi \leq 2\sin^{-1}\frac{\|u - v\|}{(1-\alpha)f(\tilde{u})} \leq 2\sin^{-1}\frac{3\,\text{width}(H)}{0.9(1-\alpha)f(\tilde{s})} < 2\sin^{-1}\frac{4\,\text{width}(H)}{f(\tilde{s})}.$$

Since $4\,\text{width}(H)/f(\tilde{s}) \leq 4W_s/f(\tilde{s})$ which is at most 0.4 by Lemma 5.1(i), we can apply (17) to conclude that $\xi < 9\,\text{width}(H)/f(\tilde{s}) \leq 9\,\text{width}(H)$. This proves (ii).

Finally, by (19), we can invoke Lemma 3.2(ii) to bound the acute angle between uv and the tangent at u. This angle is at most $\sin^{-1}\frac{\|u-v\|}{2(1-\alpha)f(\tilde{u})}$ which is less than $\xi/2$. ◻

5.1.2 Convergence lemmas

We apply the utility lemmas in the previous subsection to show that our algorithm aligns *refined*(s) quite well with the normal direction at $\tilde{s}$.

Lemma 5.3 *Assume that $\delta \leq 1/(25\rho^2)$ and $\rho \geq 5$. Let s be a sample. Let $W_s = \text{width}(\textit{refined}(s))$. For sufficiently large n, when the height of rectangle(s) is minimized, $|\theta_s| \leq 60W_s$ with probability at least $1 - O(n^{\Omega(\ln^\omega n/f_{\max})})$.*

Proof. We rotate the plane such that *refined*(s) is vertical. Suppose that $|\theta_s| > 60W_s$. We first assume that Lemmas 4.2, 4.3, 4.4, 5.1, and 5.2 hold deterministically and show that a contradiction arises with probability at least $1-O(n^{\Omega(\ln^\omega n/f_{\max})})$. Since Lemmas 4.2, 4.3, 4.4, 5.1, and 5.2 hold with probability at least $1-O(n^{\Omega(\ln^\omega n/f_{\max})})$, we can then conclude that $|\theta_s| > 60W_s$ occurs with probability at most $O(n^{\Omega(\ln^\omega n/f_{\max})})$.

Without loss of generality, we assume that $\theta_s > 0$. That is, the upward normal at s points to the left. Let L be the left boundary line of *refined*(s). By Lemma 5.1(iii), L intersects $F_\delta^- \cap$ *coarse*(s) exactly once. We use p to denote the point $L \cap F_\delta^- \cap$ *coarse*(s). We first prove the following claim which will be useful later.

Claim 1 *Orient space such that refined(s) is vertical. If $\theta_s > 60W_s$, then for any α, F_α rises strictly from left to right.*

Proof. Take any point $z \in F_\alpha \cap$ *refined*(s). By Lemma 5.1(iv), $\gamma_z \geq 0.8\theta_s - 3W_s$, which is positive as $\theta \geq 60W_s$ by assumption. Therefore, the upward normal at z points to the left, so the slope of the tangent to F_α at z is positive. ◻

Let h be the constant in Lemma 4.2. Let $k = h/1008$. Let H_1 be the strip inside *refined*(s) such that H_1 is bounded by L on the left and width$(H_1) = W_s/3$. Let H be the strip inside H_1 that is bounded by L on the left and has width $28\lambda_k\sqrt{f(\tilde{s})}$. Refer to Figure 15. Since

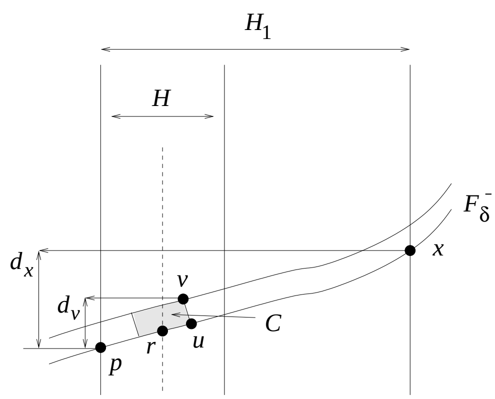

Figure 15:

$W_s \geq$ radius(*initial*(s)) which is at least $\lambda_h\sqrt{f(\tilde{s})}/3$ by Lemma 4.2,

$$\text{width}(H) = \frac{\lambda_h\sqrt{f(\tilde{s})}}{36} \leq \frac{W_s}{12}. \tag{20}$$

Thus, H lies inside H_1. Take any $(\lambda_k/\sqrt{f_{\max}})$-grid. By Lemma 5.1(iii), F_δ^- crosses H completely. Let r be the intersection point between F_δ^- and the center line of H. Let C be the $(\lambda_k/\sqrt{f_{\max}})$-cell that contains r. The distance from r to the boundary of H is $14\lambda_k\sqrt{f(\tilde{s})}$. By Lemma 3.5, the diameter of C is at most $11\lambda_k f(\tilde{r})/\sqrt{f_{\max}} \leq 11\lambda_k\sqrt{f(\tilde{r})}$. Since $f(\tilde{r}) \leq 1.1f(\tilde{s})$ by Lemma 4.5(v), the diameter of C is less than $12\lambda_k\sqrt{f(\tilde{s})}$. It follows that C lies inside H.

Let u be the rightmost vertex of C on F_δ^-. Let v be the vertex of C different from u on the normal segment at u. Let x be the intersection point between F_δ^- and the right boundary line of H_1. We are to prove that x lies above C. Since C is non-empty with very high probability, the lower side of *rectangle*(s) should intersect F_δ^- below x then. This will allow us to rotate *refined*(s) to reduce the height of *rectangle*(s) further, yielding the desired contradiction.

By Claim 1, v is the highest point in C and x is the highest point on $F_\delta^-(p,x)$. Let d_v and d_x be the height of v and x from p, respectively. Let ϕ be the acute angle between pu and the horizontal line through p. Since ϕ is at most the sum of γ_p and the angle between pu and the tangent at p, by Lemma 5.2(iii), we have $\phi \leq \gamma_p + 5\,\mathrm{width}(H)$. By Lemma 5.2(i), $\|p-u\| \leq 3\,\mathrm{width}(H)$. Observe that $d_v \leq \|p-u\| \cdot \sin\phi + \|u-v\|$. So $d_v < 3\phi\,\mathrm{width}(H) + 2\lambda_k\delta < 3\gamma_p\,\mathrm{width}(H) + 15\mathrm{width}(H)^2 + 2\lambda_k\delta$. By (20), we get $d_v < W_s\gamma_p/4 + 5W_s^2/48 + 2\lambda_k\delta$. We bound $2\lambda_k\delta$ as follows. Recall that $W_s = \min\{\sqrt{\mathrm{radius}(\textit{initial}(s))}, \mathrm{radius}(\textit{coarse}(s))/3\}$. If $W_s = \sqrt{\mathrm{radius}(\textit{initial}(s))}$, by Lemma 4.2, $W_s \geq \sqrt{\lambda_h f(\tilde{s})/3} \geq \sqrt{\lambda_h/3}$. So $2\lambda_k\delta < 2\lambda_k = \lambda_h/504 < 0.006W_s^2$. If $W_s = \mathrm{radius}(\textit{coarse}(s))/3$, by Lemma 4.4, $W_s \geq 2\sqrt{\rho}\delta/3$ and $W_s \geq \lambda_h f(\tilde{s})/3 \geq \lambda_h/3$. We get $\lambda_k = \lambda_h/1008 \leq W_s/336$ and $2\delta \leq 3W_s/\sqrt{\rho} \leq 3W_s/\sqrt{5}$, so $2\lambda_k\delta < 0.004W_s^2$. We conclude that

$$d_v < \frac{W_s\gamma_p}{4} + 0.2W_s^2.$$

Observe that px is parallel to the tangent at some point z on $F_\delta^-(p,x)$. By Lemma 5.2(ii), $\gamma_z \geq \gamma_p - 9W_s$. Since $d_x = (W_s/3)\cdot\tan\gamma_z$, we get

$$d_x \geq \frac{W_s\gamma_z}{3} \geq \frac{W_s\gamma_p}{3} - 3W_s^2.$$

Since $\theta_s > 60W_s$ by our assumption, Lemma 5.1(iv) implies that $\gamma_p \geq 0.8\theta_s - 3W_s > 45W_s$. Therefore, $d_x - d_v > W_s\gamma_p/12 - 3.2W^2 > 0$. It follows that x lies above C.

Since C is a $(\lambda_k/\sqrt{f_{\max}})$-cell, by Lemma 3.8(i), C contains some sample with probability at least $1 - n^{\Omega(\ln^\omega n/f_{\max})}$. Thus, the lower side of *rectangle*(s) lies below x with probability at least $1 - n^{\Omega(\ln^\omega n/f_{\max})}$. On the other hand, the lower side of *rectangle*(s) cannot lie below $F_\delta^- \cap H_1$, otherwise it could be raised to reduce the height of *rectangle*(s), contradiction. So the lower side of *rectangle*(s) intersects $F_\delta^- \cap H_1$ at some point a. See Figure 16.

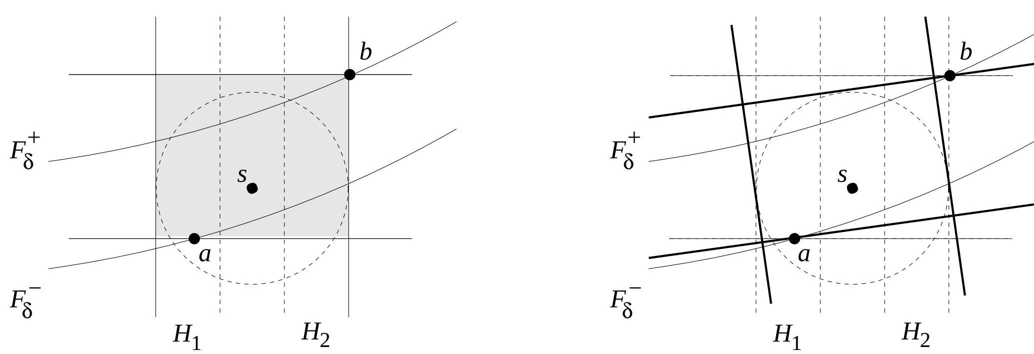

Figure 16:

Let H_2 be the strip inside *refined*(s) such that H_2 is bounded by the right boundary line of *refined*(s) on the right and $\mathrm{width}(H_2) = W_s/3$. By a symmetric argument, we can prove that the upper side of *rectangle*(s) intersects $F_\delta^+ \cap H_2$ at a point b.

As shown in Figure 16, we slightly rotate *refined*(s) in the anticlockwise direction. Since $\theta_s > 0$, the anticlockwise rotation decreases θ_s and so the rotation is legal. Moreover, as $\theta_s > 60W_s$, the small rotation keeps θ_s greater than $60W_s$. Correspondingly, we rotate the lower and upper sides of *rectangle*(s) around a and b, respectively, to obtain a rectangle R. Orient space such that the new *refined*(s) becomes vertical. By Claim 1, F_δ^- rises strictly from left to right, so F_δ^- crosses the lower side of R at most once at a from below to above. Similarly, F_δ^+ crosses the upper side of R at most once at b from below to above. This implies that R contains all the samples inside the new *refined*(s). Since a is on the left of b and below b, the anticlockwise rotation makes the width of R strictly less than the width of the old *rectangle*(s). This contradicts the assumption that the height of *rectangle*(s) is already the minimum possible. ◻

Once *refined*(s) is aligned well with the normal at $\tilde{s}$, it is intuitively true that the center point of *rectangle*(s) should lie very close to $\tilde{s}$. The following lemma proves this formally.

Lemma 5.4 *Assume that* $\delta \leq 1/(25\rho^2)$ *and* $\rho \geq 5$. *Let* s *be a sample. Let* $W_s = \text{width}(\textit{refined}(s))$. *For sufficiently large* n, *the distance between the center point of rectangle*(s) *and* $\tilde{s}$ *is at most* $(360\delta + 3)W_s$ *with probability at least* $1 - O(n^{-\Omega(\ln^\omega n/f_{\max})})$.

Proof. We first assume that Lemmas 4.2, 4.3, 4.4, 5.1, 5.2, and 5.3 hold deterministically and show that the lemma is true with probability at least $1 - O(n^{\Omega(\ln^\omega n/f_{\max})})$. Since Lemmas 4.2, 4.3, 4.4, 5.1, 5.2, and 5.3 hold with probability at least $1 - O(n^{\Omega(\ln^\omega n/f_{\max})})$, the lemma follows.

Assume that s lies on F_α^+ and the normal at $\tilde{s}$ is vertical. Let r_d (resp. r_u) be the ray that shoots downward (resp. upward) from s and makes an angle θ_s with the vertical. Let x and y be the points on F_δ^+ and F hit by r_u and r_d respectively. Let z be the point on F_δ^- hit by r_d. Let s_1 be the point on F_δ^- such that $\tilde{s_1} = \tilde{s}$. Without loss of generality, we assume that $\theta_s \geq 0$. Refer to Figure 17.

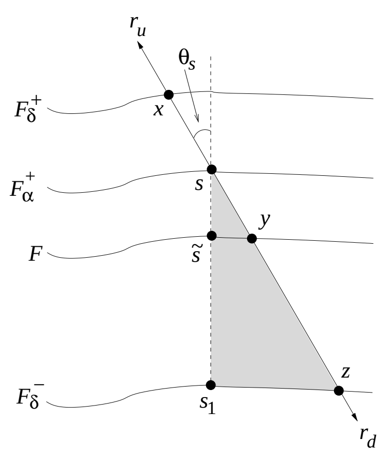

Figure 17: For the proof of Lemma 5.4.

First, we bound the distance between the midpoint of xz and y. By Lemma 4.5(iv), the acute angle between s_1z and the tangent at s_1 (the horizontal) is at most $\sin^{-1}(0.03)$. It follows that $\angle ss_1z \leq \pi/2 + \sin^{-1}(0.03)$. So $\angle szs_1 = \pi - \theta_s - \angle ss_1z \geq \pi/2 - \theta_s - \sin^{-1}(0.03)$, which is greater than 0.9 as $\theta_s \leq \pi/5$ by Lemma 5.1(ii). By applying sine law to the shaded triangle in Figure 17, we get

$$\|s_1 - z\| = \frac{\|s - s_1\| \cdot \sin\theta_s}{\sin\angle szs_1} \leq \frac{(\delta+\alpha)\theta_s}{\sin(0.9)} < 2(\delta+\alpha)\theta_s. \tag{21}$$

Similarly, we get

$$\|\tilde{s} - y\| = \frac{\|s - \tilde{s}\| \cdot \sin\theta_s}{\sin\angle sys_1} \leq \frac{\alpha\theta_s}{\sin(0.9)} < 2\alpha\theta_s. \tag{22}$$

By triangle inequality, $\|s - s_1\| - \|s_1 - z\| \leq \|s - z\| \leq \|s - s_1\| + \|s_1 - z\|$. Then (21) yields

$$(\delta+\alpha) - 2(\delta+\alpha)\theta_s \leq \|s - z\| \leq (\delta+\alpha) + 2(\delta+\alpha)\theta_s. \tag{23}$$

We can use a similar argument to show that

$$(\delta-\alpha) - 2(\delta-\alpha)\theta_s \leq \|s - x\| \leq (\delta-\alpha) + 2(\delta-\alpha)\theta_s, \tag{24}$$

$$\alpha - 2\alpha\theta_s \leq \|s - y\| \leq \alpha + 2\alpha\theta_s. \tag{25}$$

Let d_x and d_y be the distances from the midpoint of xz to x and y, respectively. Since $\|x-z\| = \|s - x\| + \|s - z\|$, by (23) and (24), we get $2\delta - 4\delta\theta_s \leq \|x - z\| \leq 2\delta + 4\delta\theta_s$. Therefore, $\delta - 2\delta\theta_s \leq d_x \leq \delta + 2\delta\theta_s$. Since $\|x - y\| = \|s - x\| + \|s - y\|$, by (24) and (25), we get $\delta - 2\delta\theta_s \leq \|x - y\| \leq \delta + 2\delta\theta_s$. We conclude that

$$d_y = |d_x - \|x - y\|| \leq 4\delta\theta_s. \tag{26}$$

Second, we bound the distance between the center point s^* of *rectangle*(s) and y. Although s^* lies on the support line of xz, it may not coincide with the midpoint of xz. There are two cases.

Case 1: the upper side of *rectangle*(s) lies above x. The upper side of *rectangle*(s) must intersect $F_\delta^+ \cap$ *refined*(s) at some point v, otherwise we could lower it to reduce the height of *rectangle*(s), contradiction. Since $\|x - v\| \leq 3W_s$ by Lemma 5.2(i), the distance between x and the upper side of *rectangle*(s) is at most $3W_s$.

Case 2: the upper side of *rectangle*(s) lies below x. Let h be the constant in Lemma 4.2. Let $k = h/84$. Take any $(\lambda_k/\sqrt{f_{\max}})$-grid. Let C be the cell that contains x.

We claim that C lies inside *refined*(s). By Lemma 3.5, the diameter of C is at most $11\lambda_k f(\tilde{x})/\sqrt{f_{\max}} \leq 11\lambda_k\sqrt{f(\tilde{x})}$. Since $f(\tilde{x}) \geq 0.9 f(\tilde{s})$ by Lemma 4.5(v), the diameter of C is less than $12\lambda_k\sqrt{f(\tilde{s})}$. Note that $W_s \geq$ radius(*initial*(s)). By Lemma 4.2, radius(*initial*(s)) $\geq \lambda_h\sqrt{f(\tilde{s})}/3 = 28\lambda_k\sqrt{f(\tilde{s})}$. So $W_s \geq 28\lambda_k\sqrt{f(\tilde{s})}$. Thus, C must lie inside *refined*(s).

Since C is a $(\lambda_k/\sqrt{f_{\max}})$-cell, by Lemma 3.8(i), C contains some sample with probability at least $1 - n^{-\Omega(\ln^\omega n/f_{\max})}$. Thus, the upper side of *rectangle*(s) cannot lie below C. It follows that the distance between x and the upper side of *rectangle*(s) is at most the diameter of C, which has been shown to be less than $W_s/2$.

Hence, the position of the upper side of *rectangle*(s) may cause s^* to be displaced from the midpoint of xz by a distance of at most $3W_s/2$. The position of the lower side of *rectangle*(s) has the same effect. So the distance between s^* and the midpoint of xz is at most $3W_s$. Since $\|s^* - y\| \leq d_y + 3W_s$, by (26), we get $\|s^* - y\| \leq 4\delta\theta_s + 3W_s$. Starting with triangle inequality, we obtain

$$\begin{aligned} \|\tilde{s} - s^*\| &\leq \|s^* - y\| + \|\tilde{s} - y\| \\ &\leq 4\delta\theta_s + 3W_s + \|\tilde{s} - y\| \\ &\stackrel{(22)}{\leq} 6\delta\theta_s + 3W_s. \end{aligned}$$

Since $\theta_s \leq 60W_s$ by Lemma 5.3, we conclude that $\|\tilde{s} - s^*\| \leq (360\delta + 3)W_s$. ◻

5.2 Homeomorphism

In this section, we prove that the output curve of the NN-crust algorithm is homeomorphic to the underlying smooth closed curve.

For each sample s, we use s^* to denote the center point of *rectangle*(s). We briefly review the processing of the center points. We first sort the center points in decreasing order of the widths of their corresponding refined neighborhoods. Then we scan the sorted list to select a subset of center points. When the current center point s^* is selected, we delete all center points p^* from the sorted list such that $\|p^* - s^*\| \leq \text{width}(\textit{refined}(s))^{1/3}$.

In the end, we call two selected center points s^* and t^* *adjacent* if $F(\tilde{s}, \tilde{t})$ or $F(\tilde{t}, \tilde{s})$ does not contain $\tilde{u}$ for any other selected center point u^*. We use G to denote the polygonal curve that connects adjacent selected center points. Clearly, if we connect $\tilde{s}$ and $\tilde{t}$ for every pair of adjacent selected center points s^* and t^*, we obtain a polygonal curve G' that is homeomorphic to the underlying smooth closed curve. Our goal is to show that the output curve of the NN-crust algorithm is exactly G. Since there is a bijection between G and G', the homeomorphism result follows.

We need to establish several technical lemmas (Lemma 5.5–5.10) before proving the homeomorphism results (Lemma 5.11 and Corollary 5.1). Throughout this section, we assume the width of any refined neighborhood is less than 1, which is true for sufficiently large n.

We first relate the widths of refined neighborhoods for two nearby center points (not necessarily selected).

Lemma 5.5 *Let p^* and q^* be two center points. If $\|\tilde{p} - \tilde{q}\| \leq f(\tilde{p})/2$, there exists a constant $\mu_1 > 0$ such that $W_q \leq \mu_1 f(\tilde{p})\sqrt{W_p}$ with probability at least $1 - O(n^{-\Omega(\ln^\omega n / f_{\max})})$.*

Proof. We prove the lemma by assuming that Lemma 4.2, 4.3, and 4.4 hold deterministically. The probability bound then follows from the probability bounds in these lemmas. For $i = p$ or q, let $R_i = \text{radius}(\textit{coarse}(i))$ and let $r_i = \text{radius}(\textit{initial}(i))$. The Lipschitz condition implies that $f(\tilde{p})/2 \leq f(\tilde{q}) \leq 3f(\tilde{p})/2$. Let h and m be the constants in Lemma 4.2.

Suppose that $W_p = \sqrt{r_p}$. By Lemma 4.2, we have

$$W_p = \sqrt{r_p} \geq \sqrt{\frac{\lambda_h \sqrt{f(\tilde{p})}}{3}} \geq \sqrt{\frac{\lambda_h}{3}} \left(\frac{2f(\tilde{q})}{3}\right)^{1/4} = \sqrt{\frac{h\lambda_m}{3m}} \left(\frac{2f(\tilde{q})}{3}\right)^{1/4}.$$

Note that $W_q \leq \sqrt{r_q}$ and $r_q \leq \sqrt{11\lambda_m f(\tilde{q})}$ by Lemma 4.2. So we get

$$W_p \geq \sqrt{\frac{h}{33m}} \left(\frac{2}{3f(\tilde{q})}\right)^{1/4} r_q \geq \sqrt{\frac{h}{33m}} \left(\frac{2}{3}\right)^{1/4} W_q^2.$$

Suppose that $W_p = R_p/3$. First, since $R_p \geq 2\sqrt{\rho}\delta$ by Lemma 4.4, we get $\rho\delta \leq 3\sqrt{\rho}W_p/2$. Second, by Lemma 4.2, $W_p \geq r_p \geq \lambda_h \sqrt{f(\tilde{p})}/3$, so we get $\sqrt{\lambda_m f(\tilde{p})} = \sqrt{m\lambda_h f(\tilde{p})/h} \leq \sqrt{3mW_p/h} \cdot f(\tilde{p})^{1/4} \leq \sqrt{3mW_p/h} \cdot f(\tilde{p})$. Finally, since $W_q \leq R_q/3$, by Lemma 4.3, we get

$$\begin{aligned} W_q &\leq \frac{5\rho\delta}{3} + \frac{\sqrt{11\lambda_m f(\tilde{q})}}{3} \\ &\leq \frac{5\rho\delta}{3} + \sqrt{\frac{11\lambda_m f(\tilde{p})}{6}} \\ &\leq \frac{5\sqrt{\rho}W_p}{2} + \sqrt{\frac{11mW_p}{2h}} \cdot f(\tilde{p}). \end{aligned}$$

□

The next result shows that the selected center points cannot be too close to each other.

Lemma 5.6 *Let p^* and q^* be two selected center points. Then $\|p^* - q^*\| > \max\{W_p^{1/3}, W_q^{1/3}\}$.*

Proof. Assume without loss of generality that p^* was selected before q^*. Since q^* was selected subsequently, q^* was not eliminated by the selection of p^*. Thus, $\|p^* - q^*\| > W_p^{1/3} \geq W_q^{1/3}$. □

Next, we bound the angle between x^*y^* and $\tilde{x}\tilde{y}$ and the angle $\angle x^*y^*z^*$ for three center points x^*, y^*, and z^*.

Lemma 5.7 *Let x^* and y^* be two center points such that $\|\tilde{x} - \tilde{y}\| \leq f(\tilde{y})/2$ and $\|x^* - y^*\| \geq W_y^{1/3}$. Then the acute angle between x^*y^* and $\tilde{x}\tilde{y}$ tends to zero as n tends to ∞ with probability at least $1 - O(n^{-\Omega(\ln^\omega n/f_{\max})})$.*

Proof. We prove the lemma by assuming that Lemmas 5.4 and 5.5 hold deterministically. The probability bound then follows from the probability bounds in these lemmas.

We translate x^*y^* to align y^* with $\tilde{y}$ and measure the acute angle θ between x^*y^* and $\tilde{x}\tilde{y}$. Let d be the distance between $\tilde{x}$ and the point $x^* + \tilde{y} - y^*$. Let $k = 360\delta + 3$. By triangle inequality and Lemma 5.4, $d \leq \|x^* - \tilde{x}\| + \|y^* - \tilde{y}\| \leq kW_x + kW_y$. Since $\|\tilde{x} - \tilde{y}\| \leq f(\tilde{y})/2$, by Lemma 5.5, $W_x \leq \mu_1 f(\tilde{y})\sqrt{W_y}$. So $d \leq k\mu_1 f(\tilde{y})\sqrt{W_y} + kW_y$. This upper bound on d is smaller than $W_y^{1/3} \leq \|x^* - y^*\|$ for sufficiently large n. So $\tilde{y}$ is further away from $x^* + \tilde{y} - y^*$ than $\tilde{x}$. It follows that θ is acute. Since d is an upper bound on the height of $x^* + \tilde{y} - y^*$ from $\tilde{x}\tilde{y}$, we have $\theta \leq \sin^{-1} \frac{d}{\|x^*-y^*\|} \leq \sin^{-1}(k\mu_1 f(\tilde{y})W_y^{1/6} + kW_y^{2/3})$. We conclude that θ tends to zero as n tends to ∞. □

Lemma 5.8 *Let x^*, y^*, and z^* be three center points such that $\tilde{y} \in F(\tilde{x}, \tilde{z})$, $\|\tilde{x} - \tilde{z}\| \leq \max\{f(\tilde{x})/4, f(\tilde{z})/4\}$, $\|x^* - y^*\| \geq W_y^{1/3}$, and $\|y^* - z^*\| \geq W_y^{1/3}$. For sufficiently large n, the angle $\angle x^* y^* z^*$ is obtuse with probability at least $1 - O(n^{-\Omega(\ln^\omega n / f_{\max})})$.*

Proof. We first show that $\|\tilde{x} - \tilde{z}\| \leq \min\{f(\tilde{x})/3, f(\tilde{z})/3\}$. Assume that $\|\tilde{x} - \tilde{z}\| \leq f(\tilde{x})/4$. By the Lipschitz condition, we have $f(\tilde{z}) \geq 3f(\tilde{x})/4$. Therefore, $\|\tilde{x} - \tilde{z}\| \leq f(\tilde{x})/4 \leq f(\tilde{z})/3$.

Let D be the disk centered at $\tilde{x}$ with radius $f(\tilde{x})/3$. Observe that $F(\tilde{x}, \tilde{z})$ lies completely inside D. Otherwise, the medial axis of F intersects the interior of D which implies that $f(\tilde{x}) \leq f(\tilde{x})/3$, contradiction. So $\|\tilde{x} - \tilde{y}\| \leq f(\tilde{x})/3$. The Lipschitz condition implies that $f(\tilde{y}) \geq 2f(\tilde{x})/3$.

Consider the angle $\angle \tilde{x}\tilde{y}\tilde{z}$. The line segments $\tilde{x}\tilde{y}$ and $\tilde{y}\tilde{z}$ are parallel to the tangents at some points on $F(\tilde{x}, \tilde{y})$ and $F(\tilde{y}, \tilde{z})$, respectively. Lemma 3.3 implies that $\angle \tilde{x}\tilde{y}\tilde{z} \geq \pi - 4\sin^{-1} \frac{\text{radius}(D)}{f(\tilde{x})} = \pi - 4\sin^{-1}(1/3) > 5\pi/9$. Since $\|\tilde{x} - \tilde{y}\| \leq f(\tilde{x})/3 \leq f(\tilde{y})/2$, by Lemma 5.7, the angle between $x^* y^*$ and $\tilde{x}\tilde{y}$ tends to zero as n tends to ∞ with probability at least $1 - O(n^{-\Omega(\ln^\omega n / f_{\max})})$. A symmetric argument shows that the angle between $y^* z^*$ and $\tilde{y}\tilde{z}$ tends to zero with probability at least $1 - O(n^{-\Omega(\ln^\omega n / f_{\max})})$ as n tends to ∞. This proves the lemma. ◻

The next lemma provides an upper bound on the the edge lengths in G.

Lemma 5.9 *Let e be an edge in G connecting two adjacent selected center points p^* and q^*. For sufficiently large n, there exists a constant $\mu_2 > 0$ such that* $\text{length}(e) \leq \mu_2 f(\tilde{p}) W_p^{1/3} + \mu_2 f(\tilde{q}) W_q^{1/3}$ *with probability at least* $1 - O(n^{-\Omega(\ln^\omega n / f_{\max})})$.

Proof. Let $k = 360\delta + 3$. Let D_p be the disk centered at p^* with radius $(1 + k\mu_1 f(\tilde{p}))W_p^{1/3}$. Let D_q be the disk centered at q^* with radius $(1 + k\mu_1 f(\tilde{q}))W_q^{1/3}$.

If D_p intersects D_q, then $\|p^* - q^*\| \leq (1 + \mu_1 f(\tilde{p}))W_p^{1/3} + (1 + \mu_1 f(\tilde{q}))W_q^{1/3}$ and we are done. Suppose that D_p does not intersect D_q. We claim that $F(\tilde{p}, \tilde{q}) \cap D_p$ is connected. Otherwise, the medial axis of F intersects the interior of D_p which implies that $f(\tilde{p}) \leq \text{radius}(D_p)$ which is less than $f(\tilde{p})$ for sufficiently large n, contradiction. Similarly, $F(\tilde{p}, \tilde{q}) \cap D_q$ is connected. It follows that $F(\tilde{p}, \tilde{q}) - (D_p \cup D_q)$ is also connected. There are two cases.

Case 1: $F(\tilde{p}, \tilde{q}) - (D_p \cup D_q)$ does not contain $\tilde{u}$ for any sample u. Let h be the constant in Lemma 4.2. Take a λ_h-partition. Since $F(\tilde{p}, \tilde{q}) - (D_p \cup D_q)$ does not contain $\tilde{u}$ for any sample u, by Lemma 3.8(i), $F(\tilde{p}, \tilde{q}) - (D_p \cup D_q)$ does not contain $F(c_i, c_{i+1})$ for any two consecutive cut-points c_i and c_{i+1} in the λ_h-partition with probability at least $1 - O(n^{-\Omega(\ln^\omega n)})$. Let y be the endpoint of $F(\tilde{p}, \tilde{q}) - (D_p \cup D_q)$ that lies on D_p. It follows that

$$|F(\tilde{p}, \tilde{q}) - (D_p \cup D_q)| < 2\lambda_h^2 f(y). \tag{27}$$

Since $\|\tilde{p} - y\| \leq 2\,\text{radius}(D_p) = 2(1 + k\mu_1 f(\tilde{p}))W_p^{1/3}$, $\|\tilde{p} - y\| \leq f(\tilde{p})/2$ for sufficiently large n. Thus, $f(y) \leq 3f(\tilde{p})/2$, so $2\lambda_h^2 f(y) < 3\lambda_h^2 f(\tilde{p})$. By Lemma 4.2, $W_p \geq \text{radius}(\mathit{initial}(p)) \geq \lambda_h \sqrt{f(\tilde{p})}/3$. So $2\lambda_h^2 f(\tilde{y}) \leq 27 W_p^2$. Substituting into (27), we get

$$|F(\tilde{p}, \tilde{q})| \leq 27 W_p^2 + 2\,\text{radius}(D_p) + 2\,\text{radius}(D_q).$$

By Lemma 5.4, $\|\tilde{p} - p^*\| \leq kW_p$ and $\|\tilde{q} - q^*\| \leq kW_q$ with probability at least $1 - O(n^{-\Omega(\ln^\omega n/f_{\max})})$. We conclude that $\|p^* - q^*\| \leq \|\tilde{p} - p^*\| + |F(\tilde{p}, \tilde{q})| + \|\tilde{q} - q^*\| \leq \mu_2 f(\tilde{p}) W_p^{1/3} + \mu_2 f(\tilde{q}) W_q^{1/3}$ for some constant $\mu_2 > 0$.

Case 2: $F(\tilde{p}, \tilde{q}) - (D_p \cup D_q)$ contains $\tilde{u}$ for some sample u. We show that this case is impossible if Lemmas 5.5 and 5.8 hold deterministically. It follows that case 2 occurs with probability at most $O(n^{-\Omega(\ln^\omega n/f_{\max})})$. We first claim that $\|p^* - u^*\| > W_p^{1/3}$. If not, Lemma 5.5 implies that $W_u \leq \mu_1 f(\tilde{p})\sqrt{W_p}$ for sufficiently large n. But then $\|p^* - \tilde{u}\| \leq \|p^* - u^*\| + \|\tilde{u} - u^*\| \leq W_p^{1/3} + kW_u \leq W_p^{1/3} + k\mu_1 f(\tilde{p})\sqrt{W_p}$. This is a contradiction as $\tilde{u}$ lies outside D_p. Similarly, $\|q^* - u^*\| > W_q^{1/3}$. So u^* is not eliminated by the selection of p^* and q^*.

Next, take any selected center point z^* different from p^* and q^* such that $\tilde{q} \in F(\tilde{u}, \tilde{z})$. We show that u^* is not eliminated by the selection of z^*. Assume to the contrary that this is false. So $\|u^* - z^*\| \leq W_z^{1/3}$. By Lemma 5.5, $W_u \leq \mu_1 f(\tilde{z})\sqrt{W_z}$ for sufficiently large n. Let $k' = 1 + k + k\mu_1$. Then $\|\tilde{u} - \tilde{z}\| \leq \|u^* - z^*\| + \|z^* - \tilde{z}\| + \|u^* - \tilde{u}\| \leq W_z^{1/3} + kW_z + kW_u \leq W_z^{1/3} + kW_z + k\mu_1 f(\tilde{z})\sqrt{W_z} \leq k' f(\tilde{z}) W_z^{1/3}$. For sufficiently large n, $k' f(\tilde{z}) W_z^{1/3} \leq f(\tilde{z})/4$. By Lemma 5.8, the angle $\angle u^* q^* z^*$ is obtuse. It follows that $\|q^* - z^*\| < \|u^* - z^*\| \leq W_z^{1/3}$, contradicting Lemma 5.6.

Symmetrically, we can show that u^* is not eliminated by any selected center point z^* different from p^* and q^* such that $\tilde{p} \in F(\tilde{z}, \tilde{u})$. In all, our algorithm should select another center point u^* such that $\tilde{u} \in F(\tilde{p}, \tilde{q}) - (D_p \cup D_q)$. This contradicts the assumption that p^* and q^* are adjacent selected center points.

◻

We are ready to show that the output curve of the NN-crust algorithm is exactly G. This will allow us to show that the output curve is homeomorphic to the underlying smooth closed curve.

Lemma 5.10 *Let p^* and q^* be two selected center points that are not adjacent. For sufficiently large n, if $\|p^* - q^*\| \leq f(\tilde{p})/4$, there is an edge e in G incident to p^* such that the angle between e and p^*q^* is acute and* $\text{length}(e) < \|p^* - q^*\|$ *with probability at least* $1 - O(n^{-\Omega(\ln^\omega n/f_{\max})})$.

Proof. Since p^* and q^* are not adjacent, there is some selected center point u^* adjacent to p^* such that $\tilde{u}$ lies on $F(\tilde{p}, \tilde{q})$ or $F(\tilde{q}, \tilde{p})$, say $F(\tilde{p}, \tilde{q})$. By Lemma 5.6, $\|p^* - u^*\| > W_u^{1/3}$ and $\|q^* - u^*\| > W_u^{1/3}$. By Lemma 5.8, the angle $\angle p^* u^* q^*$ is obtuse with probability at least $1 - O(n^{-\Omega(\ln^\omega n/f_{\max})})$. It follows that the angle between p^*u^* and p^*q^* is acute and $\|p^* - u^*\| < \|p^* - q^*\|$. ◻

Lemma 5.11 *For sufficiently large n, the output curve obtained by running the NN-crust algorithm on the selected center points is exactly G with probability at least* $1 - O(n^{-\Omega(\frac{\ln^\omega n}{f_{\max}} - 1)})$.

Proof. We first prove the lemma assuming that Lemmas 5.4, 5.8, 5.9, and 5.10 hold deterministically. We will discuss the probability bound later.

Let p^* be a selected center point. Let p^*u^* and p^*v^* be the edges of G incident to p^*. Without loss of generality, we assume that $\tilde{p}$ lies on $F(\tilde{u}, \tilde{v})$. By Lemma 5.6, $\|p^* - u^*\| > W_p^{1/3}$ and $\|p^* - v^*\| > W_p^{1/3}$.

By Lemmas 5.4 and 5.9, $\|\tilde{p} - \tilde{u}\| \leq \|\tilde{p} - p^*\| + \|\tilde{u} - u^*\| + \|p^* - u^*\| \leq kW_p + kW_u + \mu_2 f(\tilde{p})W_p^{1/3} + \mu_2 f(\tilde{u})W_u^{1/3}$, which is less than $(f(\tilde{p}) + f(\tilde{u}))/30$ for sufficiently large n. The Lipschitz condition implies that

$$0.9f(\tilde{p}) < f(\tilde{u}) < 1.1f(\tilde{p}).$$

So we get

$$\|\tilde{p} - \tilde{u}\| \leq \frac{f(\tilde{p}) + f(\tilde{u})}{30} < 0.1f(\tilde{p}), \qquad \|p^* - u^*\| \leq \frac{f(\tilde{p}) + f(\tilde{u})}{30} < 0.1f(\tilde{p}).$$

Similarly, we can show that

$$\|\tilde{p} - \tilde{v}\| < 0.1f(\tilde{p}), \qquad \|p^* - v^*\| < 0.1f(\tilde{p}).$$

Let p^*q^* be an edge computed by the NN-crust algorithm when it processes the vertex p^*. Assume to the contrary that p^*q^* is not an edge in G. If p^*q^* is computed in step 1 of the NN-crust algorithm, then q^* is the nearest neighbor of p^*. So $\|p^* - q^*\| \leq \|p^* - u^*\| < 0.1f(\tilde{p})$. By Lemma 5.10, there is another edge e in G such that length$(e) < \|p^* - q^*\|$, contradiction. Suppose that p^*q^* is computed in step 2 of the NN-crust algorithm. As we have just shown, the step 1 of the NN-crust algorithm already outputs an edge, say p^*u^*, of G where u^* is the nearest neighbor of p^*. Observe that $\|\tilde{u} - \tilde{v}\| \leq \|\tilde{p} - \tilde{u}\| + \|\tilde{p} - \tilde{v}\| < 0.2f(\tilde{p}) < 0.25f(\tilde{u})$. By Lemma 5.8, $\angle u^*p^*v^*$ is obtuse. By the NN-crust algorithm, $\angle u^*p^*q^*$ is also obtuse. Since the NN-crust algorithm prefers p^*q^* to p^*v^*, $\|p^* - q^*\| \leq \|p^* - v^*\| < 0.1f(\tilde{p})$. By Lemma 5.10, G has an edge incident to p^* that is shorter than p^*q^* and makes an acute angle with p^*q^*, contradiction.

We have shown that each output edge belongs to G. Since the NN-crust algorithm guarantees that each vertex in the output curve has degree at least two, the output curve and G have the same number of edges. So the output curve is exactly G.

Since Lemmas 5.4, 5.8, 5.9, and 5.10 hold with probability at least $1 - O(n^{-\Omega(\ln^\omega n/f_{\max})})$, the output edges incident to p^* are edges of G with probability at least $1 - O(n^{-\Omega(\ln^\omega n/f_{\max})})$. Since there are $O(n)$ output vertices, the probability that this holds for all vertices is at least $1 - O(n^{-\Omega(\frac{\ln^\omega n}{f_{\max}} - 1)})$. ◻

Corollary 5.1 *For sufficiently large n, the output curve obtained by running the NN-crust algorithm on the selected center points is homeomorphic to the underlying smooth closed curve with probability at least $1 - O(n^{-\Omega(\frac{\ln^\omega n}{f_{\max}} - 1)})$.*

Proof. We have shown that that the output curve is G. Let G' be the curve obtained by connecting $\tilde{p}$ and $\tilde{q}$ for each edge p^*q^* of G. G' is homeomorphic to the underlying smooth closed curve as p^* and q^* are adjacent. Clearly, G is homeomorphic to G' as there is a bijection between the edges of G and G'. ◻

5.3 Main theorem

We make use of the results in the previous subsections to prove the main theorem in this paper.

Theorem 5.1 *Assume that $\delta \leq 1/(25\rho^2)$ and $\rho \geq 5$. Given n noisy samples from a smooth closed curve, when n is sufficiently large, our algorithm computes a polygonal curve that satisfies the following properties with probability at least $1 - O(n^{-\Omega(\frac{\ln^\omega n}{f_{\max}} - 1)})$:*

- *Each output vertex s^* converges to $\tilde{s}$.*
- *For each output edge r^*s^*, its slope converges to the slope of the tangent at $\tilde{s}$.*
- *The output curve is homeomorphic to the smooth closed curve.*

Proof. By Lemma 5.4, an output vertex s^* converges to $\tilde{s}$ with probability at least $1 - O(n^{-\Omega(\ln^\omega n/f_{\max})})$. Since there are $O(n)$ output vertices, the pointwise convergence occurs with probability at least $1 - O(n^{-\Omega(\frac{\ln^\omega n}{f_{\max}} - 1)})$. Next, we analyze the angular differences between the tangents of the smooth closed curve and the output curve.

Let r^*s^* be an output edge. By Lemma 5.9, with probability at least $1 - O(n^{-\Omega(\ln^\omega n/f_{\max})})$, we have

$$\|r^* - s^*\| \leq \mu_2 f(\tilde{r}) W_r^{1/3} + \mu_2 f(\tilde{s}) W_s^{1/3}. \tag{28}$$

Using the above, the triangle inequality, and Lemma 5.4, we get

$$\begin{aligned} \|\tilde{r} - \tilde{s}\| &\leq \|\tilde{r} - r^*\| + \|\tilde{s} - s^*\| + \|r^* - s^*\| & (29)\\ &\leq kW_r + kW_s + \mu_2 f(\tilde{r}) W_r^{1/3} + \mu_2 f(\tilde{s}) W_s^{1/3}. & (30) \end{aligned}$$

By (28), $\|r^* - s^*\| < f(\tilde{r})/5 + f(\tilde{s})/5$ for sufficiently large n. The Lipschitz condition implies that $f(\tilde{r}) < 1.5 f(\tilde{s})$. So $\|r^* - s^*\| < f(\tilde{s})/2$. Thus, Lemma 5.5 applies and yields $W_r \leq \mu_1 f(\tilde{s}) \sqrt{W_s}$ with probability at least $1 - O(n^{-\Omega(\ln^\omega n/f_{\max})})$. Substituting into (30), we conclude that

$$\|\tilde{r} - \tilde{s}\| \leq \mu_3 f(\tilde{s})^{4/3} W_s^{1/6}, \tag{31}$$

for some constant $\mu_3 > 0$.

Let θ be the angle between $\tilde{r}\tilde{s}$ and the tangent at $\tilde{s}$. By Lemma 3.2(ii), we have $\theta \leq \sin^{-1} \frac{\mu_3 f(\tilde{s})^{1/3} W_s^{1/6}}{2}$. Let θ' be the acute angle between r^*s^* and $\tilde{r}\tilde{s}$. By (31), $\|\tilde{r} - \tilde{s}\| \leq f(\tilde{s})/2$ for sufficiently large n. Thus, by Lemma 5.7, θ' tends to zero as n tends to ∞ with probability at least $1 - O(n^{-\Omega(\ln^\omega n/f_{\max})})$. We conclude that $\theta + \theta'$ tends to zero as n tends to ∞, so the slope of r^*s^* converges to the slope of the tangent at $\tilde{s}$. Since there are $O(n)$ output edges, the convergence of their slopes occur with probability at least $1 - O(n^{-\Omega(\frac{\ln^\omega n}{f_{\max}} - 1)})$.

The output curve is homeomorphic to the smooth closed curve by Corollary 5.1. ∎

6 Discussion

We expect that the approach will also work for handling curves with features: the sampled "curve" consists of a collection of simple curve segments that may only share endpoints, thus forming features like corners, branchings and terminals. Some previous works have already considered terminal and corner points. Allowing branchings extends this to the most general problem. Furthermore, we aim to handle features in the presence of noise. A motivation for allowing branchings is that if we consider surfaces in 3-d with features like sharp edges and corners, then these form a curve graph (in 3-d) with corners, branchings and terminals. The output reconstruction is expected to identify the features as part of the reconstruction. As in previous works, the definition of local feature size is modified to avoid a zero local feature size in corners and branchings points, by pruning the medial axis near the features. The shape fitting can be done by finding a branching of k slabs – the Minkowski sum of a disk and k rays originating from a common apex (see figure) – with smallest width that contains the points. Almost brute force algorithms for these fitting problems run in polynomial time. Linear time approximation algorithms seem possible by adapting recent work on k-line centers [1].

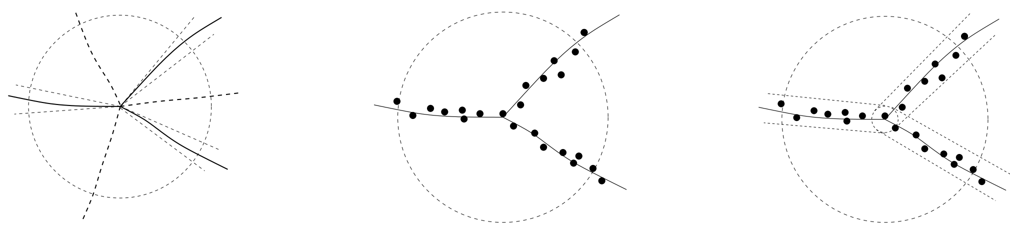

Figure 18: Degree 3 branching, Noisy sampling and Fitting.

We also need a Modified NN-Crust that works correctly for a noise free locally uniform sampling from a curve with features. Such a variant is possible if we assume that for each feature in the curve, the sampling should include a sample s which is identified and provided with a k cones corresponding to the incident curve branches. This is the case for us, since this is information is obtained from the feature fitting step. In the Modified NN-Crust, each feature sample s selects the nearest neighbor in each of its cones, then each non-feature sample s that was not selected by a feature sample proceeds as in the NN-Crust, and each non-feature s that was selected by a feature sample s' selects the nearest neighbor in a cone opposite to s'

To guarantee that the original curve is reconstructed, a very restricted (locally uniform) sampling condition is needed: as it has been pointed out before, the sampling can "simulate" non-existing features and "destroy" real ones. So, a witness guarantee as in [5] is desirable. Beyond this, we also use uniformity of the sampling to assure that the type of the neighborhood can be determined locally. To avoid this, the steps of neighborhood identification and global reconstruction should be interconnected. For example, though at a small scale, a neighborhood may seem to contain a terminal, it may be that this is not the case and that this is only realized when a global consistent reconstruction is not possible under this assumption. Appropriate rules for the interaction between feature fitting and reconstruction need to be explored.

An integration of fitting and reconstruction is also necessary to avoid our current assumption of dense noise. In a different direction, it seems possible to handle outliers if the algorithm uses shape fitting with outliers.

References

[1] P. K. Agarwal, C. M. Procopiuc and K. Varadarajan. Approximation algorithms for k-line center. *Proc. 10th Annu. European Sympos. Alg.,* 2002.

[2] E. Althaus and K. Mehlhorn. Traveling salesman based curve reconstruction in polynomial time. *SIAM J. Comput.*, 31 (2001), 27–66.

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

[4] T. K. Dey and P. Kumar. A simple provable algorithm for curve reconstruction. *Proc. 10th. Annu. ACM-SIAM Sympos. Discrete Alg.*, 1999, 893–894.

[5] T. K. Dey, K. Mehlhorn and E. Ramos. Curve reconstruction: connecting dots with good reason. *Comput. Geom. Theory & Appl.*, 15 (2000), 229–244.

[6] T. K. Dey and R. Wenger. Reconstructing curves with sharp corners. *Comput. Geom. Theory & Appl.*, 19 (2001), 89–99.

[7] T. K. Dey and R. Wenger. Fast reconstruction of curves with sharp corners. *Comput. Geom. Theory & Appl.*, 12 (2002), 353–400.

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

[9] S. Funke and E. A. Ramos. Reconstructing a collection of curves with corners and end-points. *Proc. 12th Annu. ACM-SIAM Sympos. Discrete Alg.*, 2001, 344–353.

[10] J. Giesen. Curve reconstruction, the Traveling Salesman Problem and Menger's Theorem on length. *Discrete & Comp. Geom.*, 24 (2000), 577–603.

[11] C. Gold and J. Snoeyink. A one-step crust and skeleton extraction algorithm. *Algorithmica*, 30 (2001), 144–163.