Tighter Bounds for Reconstruction from ϵ-Samples

Håvard Bakke Bjerkevik
Institute of Geometry, Technische Universität Graz, Austria

Abstract

We show that reconstructing a curve in $\mathbb{R}^d$ for $d \geq 2$ from a 0.66-sample is always possible using an algorithm similar to the classical NN-Crust algorithm. Previously, this was only known to be possible for 0.47-samples in $\mathbb{R}^2$ and $\frac{1}{3}$-samples in $\mathbb{R}^d$ for $d \geq 3$. In addition, we show that there is not always a unique way to reconstruct a curve from a 0.72-sample; this was previously only known for 1-samples. We also extend this non-uniqueness result to hypersurfaces in all higher dimensions.

2012 ACM Subject Classification Mathematics of computing → Geometric topology

Keywords and phrases Curve reconstruction, surface reconstruction, ϵ-sampling

Digital Object Identifier 10.4230/LIPIcs.SoCG.2022.9

Related Version *Full Version*: https://arxiv.org/abs/2112.03656v2 [8]

Funding The author is supported by the Austrian Science Fund (FWF) grant number P 33765-N.

Acknowledgements The author would like to thank Michael Kerber for insights into the size of Delaunay triangulations, and Stefan Ohrhallinger and Scott A. Mitchell for answering my questions about the state of the art of curve and surface reconstruction.

1 Introduction

The main problem considered in this paper is that of *curve reconstruction*. Given a (finite) set of points $\mathcal{S}$ in $\mathbb{R}^d$, we assume that this is a subset of a union $\mathcal{C}$ of closed curves, and we want to reconstruct $\mathcal{C}$ knowing only $\mathcal{S}$. Reconstructing $\mathcal{C}$ exactly from a finite set of points is unfeasible, so we restrict the problem to finding the graph $G_{\mathcal{C}}(\mathcal{S})$ on $\mathcal{S}$ induced by $\mathcal{C}$: there is an edge in $G_{\mathcal{C}}(\mathcal{S})$ between two points in $\mathcal{S}$ if you can walk from one to the other along $\mathcal{C}$ without meeting another point of $\mathcal{S}$.

To do this, one needs an assumption on $\mathcal{S}$ and $\mathcal{C}$. Some work on curve reconstruction and similar problems uses global assumptions for instance related to the maximum curvature [5, 7, 12, 21, 24, 25]. A weakness of this approach is that it may force you to sample the whole curve densely even if just a small portion of it has large curvature. An influential paper by Amenta, Bern and Eppstein [3] introduced the Crust algorithm along with a local sampling condition allowing the sampling density to vary depending on the local distance to the *medial axis* of $\mathcal{C}$. To be precise, they guarantee correct reconstruction for any *ϵ-sampled* curve in the plane whenever $\epsilon < 0.252$. The condition that a curve is ϵ-sampled is weaker the larger ϵ is, so we would like to guarantee correct reconstruction for ϵ-sampled curves for as large an ϵ as possible.

There followed a number of papers seeking to improve the sampling conditions of [3]: Dey and Kumar [15] introduced NN-Crust (NN = nearest neighbor), which allows curves in higher-dimensional space, and prove that correct reconstruction is guaranteed for $\epsilon < \frac{1}{3}$; Lenz [20] defines a family of algorithms of which NN-Crust is a special case and conjectures that $\epsilon \leq 0.48$ is sufficient for correctness in another special case; and Ohrhallinger et al. [22] introduce HNN-Crust, proving correct reconstruction for $\epsilon < 0.47$, and also for $\rho < 0.9$, where ρ is a *reach*-based parameter that is related to (but different from) the parameter ϵ. It is shown in [3, Observation 6] that correct reconstruction cannot be guaranteed for $\epsilon \geq 1$.

In addition, there have been several papers improving on [3] in other ways: Gold [18] simplified the Crust algorithm; Dey et al. [16] gave an algorithm allowing open curves; and Dey and Wenger [17] considered curves with corners. Finally, we mention that [2] ties the ϵ-sampling condition to a completely different approach to curve reconstruction by showing that a solution of the traveling salesman problem on the sample points gives a correct reconstruction from an ϵ-sample for $\epsilon < 0.1$. For further references, we refer to the recent survey of Ohrhallinger et al. [23] on curve reconstruction in the plane.

Moving up to higher dimensions, one can consider the problem of *submanifold reconstruction* [1, 10, 13, 14, 21]. Instead of working with samples of a curve, one assumes that the points are sampled from a submanifold in $\mathbb{R}^d$ for $d \geq 3$; the case of surfaces in $\mathbb{R}^3$ is of particular interest. While this is not the main focus of the paper, we note that this problem is important from a practical point of view; see for instance [6] for a survey covering the literature related to 3D scannings with imperfections. So far, the results using ϵ-sampling have been much weaker for surface reconstruction than for curve reconstruction. For $d = 3$, correct surface reconstruction is only known to be possible to guarantee for $\epsilon \leq 0.06$ [4].

1.1 Our contributions

The question we study is: For which ϵ is it possible to guarantee correct curve reconstruction using an ϵ-sample? Despite the popularity of ϵ-sampling as a sampling condition in the literature and the body of work aiming to weaken sampling conditions, there is still a large gap between the ϵ for which we know that reconstruction is always possible and the ϵ for which we know that it is not always possible: For any $\epsilon \in (0.47, 1)$, it is as far as the author knows an open question if it is possible to guarantee correct reconstruction of a curve (or union of curves) in $\mathbb{R}^2$ using an ϵ-sample. For curves in $\mathbb{R}^d$, $d \geq 3$, the same is true for $\epsilon \in (\frac{1}{3}, 1)$. We improve this situation drastically in both ends. First we describe algorithms that guarantee correct reconstruction for $\epsilon = 0.66$ for all $d \geq 2$. Algorithm 1 runs in $O(n^2)$ for any fixed d, and Algorithm 2 runs in $O(n \log n)$ for $d = 2$. While we have not implemented our algorithms, we believe that the speed of Algorithm 2 in practice is comparable to that of the algorithms in [15] and [22] because of their similarities.

Secondly, we give an example demonstrating that one cannot in general guarantee correct reconstruction using 0.72-samples for any $d \geq 2$. Thus, the interval of ϵ for which it is unknown if an ϵ-sample is enough for reconstruction is reduced from $(0.47, 1)$ (or $(\frac{1}{3}, 1)$ for $d \geq 3$) to $(0.66, 0.72)$.

By a straightforward generalization, we use our example to prove that a 0.72-sample is not in general enough to guarantee correct reconstruction of a manifold of any dimension. We do not show any positive results in higher dimensions, but we hope that since we do not put any restriction on the ambient dimension of the set of samples, our ideas can be useful also for reconstruction of higher-dimensional manifolds.

A serious alternative to the ϵ-sampling condition is the ρ-sampling condition of [22]. The authors of [22] argue that ϵ-sampling with $\epsilon \leq 0.47$ requires more sample points than what ρ-sampling does. With our new bounds on ϵ, the situation changes somewhat. An in-depth discussion of the relationship between ϵ-sampling and ρ-sampling is beyond the scope of this paper (as is the question of whether the two sampling conditions can be combined in a way that exploits the advantages of both of them), but we study some instructive examples in the full version of the paper [8, Appendix B]. To summarize, ρ-sampling seems to do better for curves with slowly changing curvature, while ϵ does better in some examples with rapidly changing curvature. Both our upper and lower bounds for ϵ help us understand the relative strengths of ϵ- and ρ-sampling.

We begin by introducing necessary definitions and notation in Section 2, before we prove the main theorem in Section 3. In Section 4, we show that correct reconstruction from 0.72-samples is not always possible, and we finish off by generalizing the example to higher dimensions in Section 5.

2 Definitions and notation

Throughout most of the paper, we work with a finite, disconnected union $\mathcal{C}$ of closed curves in $\mathbb{R}^d$ for some fixed $d \geq 2$, and a finite subset $\mathcal{S}$ of $\mathcal{C}$. We will call the elements of $\mathcal{S}$ *sample points*. By a closed curve, we mean the image of an injective map from the circle. Sometimes it will be convenient to fix an orientation of (a connected component of) $\mathcal{C}$. The notation $a \to b$ means that we have chosen an orientation of a connected component of $\mathcal{C}$ containing $a, b \in \mathcal{S}$ and that by starting at a and moving along $\mathcal{C}$ following this orientation, the next element of $\mathcal{S}$ one encounters is b. We use the shorthand $a \to b \to c$ when we mean $a \to b$ and $b \to c$. For p, q in the same connected component of $\mathcal{C}$, we define $[p, q]$ as $\{p\}$ if $p = q$, and as the image of any injective path from p to q that is consistent with the orientation of $\mathcal{C}$ if $p \neq q$. We define $[a, b)$, $(a, b]$ and (a, b) similarly depending on whether a and/or b are included or not. By a *midpoint of* $[a, b]$ we mean a point $p \in [a, b]$ with $d(p, a) = d(p, b)$, where $d(x, y)$ denotes Euclidean distance.

If $(a \to)b \to c$ or $c \to b(\to a)$, we say that $(a,)$ b and c are *consecutive*. We define $G_{\mathcal{C}}(\mathcal{S})$ as the graph on $\mathcal{S}$ with an edge between a and b if and only if a and b are consecutive.

For $X \subset \mathbb{R}^d$, let $d(x, X) := \inf_{y \in X} d(x, y)$. The *medial axis* $\mathcal{M}$ [9] is the set of points in $\mathbb{R}^d$ that do not have a unique closest point in $\mathcal{C}$. For $p \in \mathcal{C}$, the *local feature size* $\operatorname{lfs}(p)$ is defined as $d(p, \mathcal{M})$. For $\epsilon > 0$, we say that $\mathcal{S} \subset \mathcal{C}$ is an *ϵ-sample* (of $\mathcal{C}$) if for all $p \in \mathcal{C}$, $d(p, \mathcal{S}) < \epsilon \operatorname{lfs}(p)$. Note that being an ϵ-sample is a stronger condition the smaller ϵ is. Throughout the paper we will assume that $\mathcal{S}$ is an ϵ-sample, but our assumptions on ϵ will vary.

We define $\operatorname{cl}\colon \mathbb{R}^d \setminus \mathcal{M} \to \mathcal{C}$ by letting $\operatorname{cl}(x)$ be the point in $\mathcal{C}$ closest to x; i.e., $\operatorname{cl}(x) = \arg\min_{p \in \mathcal{C}} d(x, p)$. It follows immediately from the definition of $\mathcal{M}$ that cl is well-defined. We prove that cl is continuous in Lemma 2.

We use the notation $B_x(r)$ for the closed ball with radius r centered at $x \in \mathbb{R}^d$. For $x, y \in \mathbb{R}^d$, the closed line segment from x to y is denoted by $\overline{xy}$.

We often restrict our attention to a plane $\Pi \subset \mathbb{R}^d$, which we identify with $\mathbb{R}^2$. This way, we can associate canonical coordinates (x, y) to each point $p \in \Pi$.

3 Proof that 0.66-samples allow reconstruction

This section is devoted to giving a proof of the main theorem:

▶ **Theorem 1.** *Let $\mathcal{C}$ be a union of closed curves in $\mathbb{R}^d$ for some $d \geq 2$, and let $\mathcal{S}$ be a 0.66-sample of $\mathcal{C}$ containing n points. Given $\mathcal{S}$ as input,* NN-compatible *and* Compatible-crust *both compute $G_{\mathcal{C}}(\mathcal{S})$. The former runs in $O(n^2)$, and for $d = 2$, the latter runs in $O(n \log n)$.*

The algorithms are rather simple, and are similar to the previous Crust-type algorithms. To be specific, Compatible-crust borrows the idea from [3] of only selecting edges from the Delaunay triangulation[1], and both algorithms use the idea from [15] of including an edge

[1] For an introduction to Delaunay triangulations in the plane, see [11, Chapter 9].

between each sample point and its nearest neighbor (called "closest" in the algorithms) in addition to the nearest neighbor satisfying some condition related to the angle between the resulting two edges (called "clComp" in the algorithms). The new ingredient in our algorithm is that we require triples of consecutive points to be *compatible* (see Figure 3), which is a different criterion than those used in previous algorithms. We define this compatibility property in Section 3.3. This criterion has the advantage over criteria used in previous papers in that it is the optimal local criterion for when a triple of points can be consecutive: If a triple is not compatible, it cannot be consecutive, while if it is compatible, there is a curve passing through the three points that does not violate the sampling condition locally. It will be clear from the definition that checking if a triple $(a, b, c) \in \mathcal{S}^3$ is compatible can be done in constant time. The separation into two algorithms is done to optimize the running time: For $d = 2$, computing the Delaunay triangulation saves us time, while for $d \geq 3$, a more straightforward approach is at least as efficient in the worst case.

Algorithm 1 NN-COMPATIBLE.

Input: 0.66-sample $\mathcal{S} \subset \mathbb{R}^d$ of $\mathcal{C}$ for $d \geq 2$
Output: $G_{\mathcal{C}}(\mathcal{S})$
Initialize $G \leftarrow \{\}$
foreach $x \in \mathcal{S}$ **do**
  closest $\leftarrow \arg\min_{y \in \mathcal{S} \setminus \{x\}} \{d(x, y)\}$
  CompNeigh $\leftarrow \{y \in \mathcal{S} \mid (\text{closest}, x, y) \text{ is compatible}\}$
  clComp $\leftarrow \arg\min_{y \in \text{CompNeigh}} \{d(x, y)\}$
  $G \leftarrow G \cup \{\{x, \text{closest}\}, \{x, \text{clComp}\}\}$
return G

In NN-COMPATIBLE, we run through the for-loop n times. Each line in the loop can be executed in $O(n)$, which gives a total running time of $O(n^2)$.

Algorithm 2 COMPATIBLE-CRUST.

Input: 0.66-sample $\mathcal{S} \subset \mathbb{R}^d$ of $\mathcal{C}$ for $d \geq 2$
Output: $G_{\mathcal{C}}(\mathcal{S})$
Compute the 1-skeleton $D_1(\mathcal{S})$ of a Delaunay triangulation of $\mathcal{S}$.
Initialize $G \leftarrow \{\}$
foreach $x \in \mathcal{S}$ **do**
  Neigh $\leftarrow$ the set of vertices in $D_1(\mathcal{S})$ adjacent to x
  closest $\leftarrow \arg\min_{y \in \text{Neigh}} \{d(x, y)\}$
  CompNeigh $\leftarrow \{y \in \text{Neigh} \mid (\text{closest}, x, y) \text{ is compatible}\}$
  clComp $\leftarrow \arg\min_{y \in \text{CompNeigh}} \{d(x, y)\}$
  $G \leftarrow G \cup \{\{x, \text{closest}\}, \{x, \text{clComp}\}\}$
return G

Computing a Delaunay triangulation in the plane can be done in $O(n \log n)$ [11, Theorem 9.12]. The total number of edges in $D_1(\mathcal{S})$ is $O(n)$, so the sum of the sizes of all the Neigh over all $x \in \mathcal{S}$ is $O(n)$. Thus, the total running time of the for-loop is $O(n)$. This gives a running time for COMPATIBLE-CRUST of $O(n \log n + n) = O(n \log n)$ for $d = 2$. For $d \geq 3$, the Delaunay triangulation may have a size as large as $\Theta(n^{\lceil d/2 \rceil})$ [19, Chapter 27.1], in which case COMPATIBLE-CRUST does not do better than NN-COMPATIBLE for $d \in \{3, 4\}$ and does worse for $d \geq 5$.

It remains to be proved that the algorithms output $G_{\mathcal{C}}(\mathcal{S})$. Since Compatible-crust restricts itself to the set of edges of the Delaunay triangulation, we need to know that this set contains the edges of $G_{\mathcal{C}}(\mathcal{S})$. In the planar case, this is proved in [3, Lemma 11]. We extend the result to higher ambient dimensions in Corollary 5.

Finally, we need to prove that the closest and "closest compatible" neighbors to a sample point are indeed the adjacent vertices in $G_{\mathcal{C}}(\mathcal{S})$. As the proof is rather long and technical, we devote a full section to it, which we split into three subsections: In Section 3.1, we prove a sequence of lemmas about the local behavior of $\mathcal{S}$ and $\mathcal{C}$. Then, in Section 3.2, we prove lower bounds on the angle between certain triples of points on $\mathcal{C}$; in particular, Lemma 11 implies that consecutive triples of points have to be compatible. Lastly, in Section 3.3, we use the results from the first two subsections to prove that the edges constructed by the algorithms are indeed exactly the edges in $G_{\mathcal{C}}(\mathcal{S})$. Some of the proofs are omitted and appear only in the appendix of the full version of the paper [8].

3.1 Basic observations about $\mathcal{S}$ and $\mathcal{C}$

Recall that $\mathcal{S}$ is assumed to be an ϵ-sample of $\mathcal{C}$. In this subsection, we assume $\epsilon \leq 1$. Later, we will restrict ϵ to smaller values and state our assumptions on ϵ explicitly in each case.

For $p \in \mathcal{C}$, define $d_p = d(p, \mathcal{S})$. By definition of cl and ϵ-sample, cl is defined in $B_p\left(\frac{d_p}{\epsilon}\right)$. Since we assume $\epsilon \leq 1$, cl is in particular defined in $B_p(d_p)$. We will use the following lemma throughout the paper without referring to it explicitly.

▶ **Lemma 2.** *cl is continuous.*

Proof. Let $x \in \mathbb{R}^d \setminus \mathcal{M}$, and let $x_1, x_2, \ldots$ be a sequence of points in $\mathbb{R}^d \setminus \mathcal{M}$ that converges to x. To show that cl is continuous, it is enough to show that the image of the sequence under cl converges to $\mathrm{cl}(x)$. Let y be an accumulation point in $\mathcal{C}$ of the sequence $\mathrm{cl}(x_1), \mathrm{cl}(x_2), \ldots$, which exists by compactness of $\mathcal{C}$. Then $d(x, y) \leq d(x, y')$ for any $y' \in \mathcal{C}$, so $y = \mathrm{cl}(x)$. Thus, $\mathrm{cl}(x)$ is the only accumulation point of $\mathrm{cl}(x_1), \mathrm{cl}(x_2), \ldots$, so by compactness of $\mathcal{C}$, the sequence converges to $\mathrm{cl}(x)$. ◀

▶ **Lemma 3.** *Let $x \in \mathbb{R}^d$ and $q \in \mathcal{C}$ be such that $\overline{xq}$ does not intersect the medial axis. Let $p = cl(x)$. Then the interior of $B_x(d(x, q))$ contains either $[p, q)$ or $(q, p]$.*

Proof. By continuity of cl and connectedness of $\overline{xq}$, $\mathrm{cl}(\overline{xq})$ must contain either $[p, q]$ or $[q, p]$. Suppose the former. Then for any $z \in [p, q)$, $z = \mathrm{cl}(i)$ for some $i \in \overline{xq}$. Thus,

$$d(x, q) = d(x, i) + d(i, q) > d(x, i) + d(i, z) \geq d(x, z).$$

The statement follows, and the argument for $[q, p]$ is exactly the same. ◀

▶ **Lemma 4.** *Let $a \to b$ and $p \in (a, b)$. Then $d_p = \min\{d(p, a), d(p, b)\}$, and $d_p < d(p, s)$ for all $s \in \mathcal{S} \setminus \{a, b\}$.*

Proof. Suppose $s \notin \{a, b\}$ is a point in $\mathcal{S}$ minimizing the distance to p, so $d_p = d(p, s)$. Then $B_p(d(p, s)) = B_p(d_p)$ and thus $\overline{ps}$ does not intersect the medial axis. Since $\mathrm{cl}(p) = p$, Lemma 3 (with $x = p$ and $q = s$) shows that the interior of $B_p(d_p)$ contains either a or b, which is a contradiction, as then either $d(p, a)$ or $d(p, b)$ would be smaller than $d(p, s)$. Thus, d_p is equal to either $d(p, a)$ or $d(p, b)$. ◀

As a step in proving the correctness of Compatible-crust, we need to show that for $a \to b$, there is an edge between a and b in the Delaunay triangulation of $\mathcal{S}$. Since we do not assume that $\mathcal{S}$ is in general position, we do not know that there is a unique Delaunay

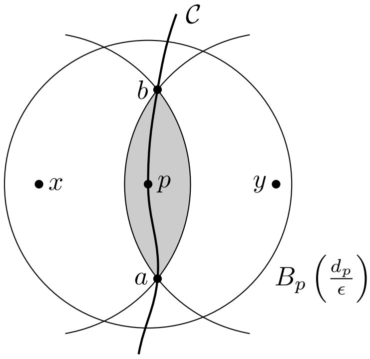

Figure 1 The planar case with $X(a,b) = \{x,y\}$. The shaded area is $U(a,b)$ and contains $[a,b]$ by Lemma 8 (iii). By Lemma 8 (ii), $B_p\left(\frac{d_p}{\epsilon}\right)$ contains $X(a,b)$, where p is the midpoint on $[a,b]$.

triangulation of $\mathcal{S}$. Still, we know that if there is a closed ball B such that $B \cap \mathcal{S} = \{a,b\}$, then any Delaunay triangulation of $\mathcal{S}$ has an edge between a and b. In the special case of curves in the plane, the following was proved in [3, Lemma 11].

▶ **Corollary 5.** *Let $a \to b$. Then there is an edge between a and b in any Delaunay triangulation of $\mathcal{S}$.*

Proof. Let p be a midpoint on $[a,b]$. By Lemma 4, $B_p(d_p) \cap \mathcal{S} = \{a,b\}$, so there is an edge between a and b in the Delaunay triangulation of $\mathcal{S}$. ◀

For $x,y \in \mathbb{R}^d$, let $E(x,y)$ be the set of points in $\mathbb{R}^d$ that are equidistant from x and y.

▶ **Lemma 6.** *Let $b \in \mathcal{S}$, let $a \neq b$ be in the same connected component of $\mathcal{C}$ as b, let p be either the midpoint on $[a,b]$ or equal to a, and assume $d_p = d(p,b)$. Then for every $x \in B_p\left(\frac{d_p}{\epsilon}\right) \cap E(a,b)$,*
(i) *$cl(x) \in (a,b)$,*
(ii) *$(a,b) \subset B_x(d(x,b))$.*

Proof. (i): Let $B = B_p\left(\frac{d_p}{\epsilon}\right)$, and let m be the midpoint on $[a,b]$. If $p = a$, then by Lemma 3, $m \in B$. Trivially, $m \in B$ also holds if $p = m$. Since $\mathcal{S}$ is an ϵ-sample, B does not intersect the medial axis, so $\mathrm{cl} : B \to \mathcal{C}$ is well-defined. Clearly, $\mathrm{cl}(m) = m$, and $a,b \notin \mathrm{cl}(B \cap E(a,b))$, as $d(a,x) = d(b,x)$ for every $x \in E(a,b)$. Since cl is continuous and $B \cap E(a,b)$ connected, we get that $\mathrm{cl}(B \cap E(a,b)) \subset (a,b)$.

(ii): Since $\overline{xb} \subset B_p\left(\frac{d_p}{\epsilon}\right)$, Lemma 3 tells us that $[\mathrm{cl}(x),b)$ is in the interior of $B_x(d(x,b))$ (since $a \in (b,\mathrm{cl}(x)]$ is not in the interior of $B_x(d(x,b)) = B_x(d(x,a))$), and so must $(a,\mathrm{cl}(x)]$ by a symmetric argument. ◀

▶ **Definition 7.** *For $a \neq b \in \mathbb{R}^d$, let $X(a,b)$ be the set of x such that $d(x,a) = d(x,b) = \frac{d(a,b)}{\epsilon\sqrt{4-\epsilon^2}}$, and let $U(a,b) = \bigcap_{x \in X(a,b)} B_x\left(\frac{d(a,b)}{\epsilon\sqrt{4-\epsilon^2}}\right)$, which is equal to $\bigcap_{x \in X(a,b)} B_x\left(d(x,a)\right)$.*

▶ **Lemma 8.** *Let $a \to b$.*
(i) *Let $p' \in E(a,b) \cap \partial U(a,b)$, and let x be the point in $X(a,b)$ maximizing the distance to p'. Then $d(p',a) = \epsilon d(p',x)$, $d(p',x) = d(a,x)$ and $2\angle axp' = \angle axb$.*
(ii) *Let p be the midpoint of $[a,b]$. Then $X(a,b) \subset B_p\left(\frac{d_p}{\epsilon}\right)$.*
(iii) *$(a,b) \subset U(a,b)$.*

See Figure 1 for an illustration of (ii) and (iii). We prove the lemma in [8, Appendix A.1].

3.2 Restrictions of angles between points on $\mathcal{C}$

With help from the results of the previous subsection, we now prove results that essentially limit the curvature of $\mathcal{C}$ locally.

▶ **Proposition 9.** *Let $\epsilon \le 0.765$, and let $a \to b \to c$ with $p \in (a,b)$ and $d(p,b) \le d(p,a)$. Then for any x such that $d(x,p) = d(x,b) = \frac{d_p}{\epsilon}$, $(b,c] \cap B_x\left(\frac{d_p}{\epsilon}\right) = \emptyset$.*

The rough idea of the proof is to assume there is a $c' \in (b,c] \cap B_x\left(\frac{d_p}{\epsilon}\right)$ and consider a line segment $\overline{xm}$, where x satisfies the conditions in the lemma and m is the midpoint on $\overline{bc'}$. One can show that cl is defined on $\overline{xm}$, that $\mathrm{cl}(x) \in (p,b)$, and that $\mathrm{cl}(m) \in (b,c')$ and derive that $\mathrm{cl}(\overline{xm})$ is disconnected, which is a contradiction by continuity of cl. We give the full details in [8, Appendix A.2].

▶ **Corollary 10.** *Let $\epsilon \le 0.66$, let $a \to b \to c$, and let p be the midpoint of $[a,b]$ and $q \in (b,c]$. Then*

$$\angle pbq > 70.73° + \arccos\left(0.33\frac{d(q,b)}{d(p,b)}\right).$$

In particular, if $d(p,b) \ge d(q,b)$, then $\angle pbq > 141°$.

Proof. We restrict our attention to a plane containing p, b and q and assume without loss of generality that $p = (0,-1)$, $b = (0,0)$ and that q is not to the left of the y-axis. By Proposition 9, q cannot be in the disc D with radius $\frac{1}{\epsilon}$ with p and b on the boundary and center x to the right of the y-axis. Under this condition, we have $\angle pbq > \angle pbq'$, where q' is on the boundary of D above the x-axis and $d(q',b) = d(q,b)$. As illustrated in Figure 2a, $\cos\angle pbx = \frac{1/2}{1/\epsilon} \le 0.33$. Similarly, $\cos\angle xbq' = \frac{d(q',b)/2}{1/\epsilon} \le 0.33d(q',b)$. Since arccos is decreasing, we get

$$\begin{aligned}\angle pbq &> \angle pbq'\\ &= \angle pbx + \angle xbq'\\ &\ge \arccos(0.33) + \arccos(0.33d(q',b))\\ &> 70.73° + \arccos(0.33d(q',b)).\end{aligned}$$

If we do not assume $d(p,b) = 1$, we have to replace $d(q',b)$ with $\frac{d(q',b)}{d(p,b)}$ in the last expression. Since $d(q',b) = d(q,b)$, this yields the wanted inequality. If $d(p,b) \ge d(q,b)$, then this lower bound is weakest when $d(q,b) = d(p,b)$. In this case the right-hand side is $> 141.46°$. ◀

▶ **Lemma 11.** *Let $\epsilon \le 0.765$, and let $a \to b \to c$. Then $(b,c] \cap B_x(d(x,a)) = \emptyset$ for all $x \in X(a,b)$.*

This proof is similar to that of Proposition 9; see [8, Appendix A.3] for the details.

▶ **Definition 12.** *We call a triple (a,b,c) of sample points* compatible *if $c \notin B_x(d(x,b))$ for all $x \in X(a,b)$ and $a \notin B_y(d(y,b))$ for all $y \in X(b,c)$.*

See Figure 3. Lemma 11 then implies that if $a \to b \to c$, then (a,b,c) is compatible.

▶ **Lemma 13.** *Let $\epsilon \le 0.66$ and suppose (a,b,c) is compatible. Then*

$$\angle abc > 51.45° + \arccos\left(0.6231\frac{d(c,b)}{d(a,b)}\right).$$

In particular, $\angle abc > 102.9°$.

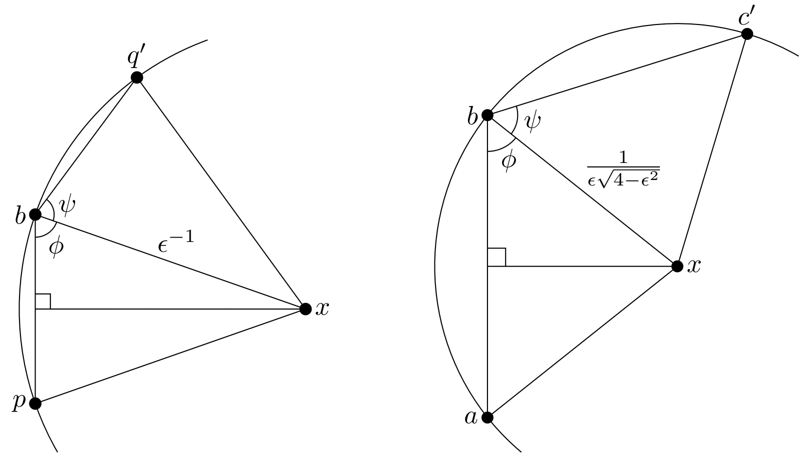

(a) If $d(p,b)=1$, then $\cos\phi = \frac{1/2}{1/\epsilon}$.

(b) If $d(a,b)=1$, then $\cos\phi = \frac{\epsilon\sqrt{4-\epsilon^2}}{2}$.

Figure 2

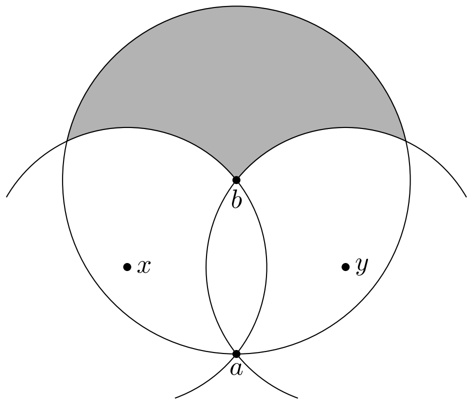

Figure 3 The planar case with $X(a,b)=\{x,y\}$. If $d(b,c)\le d(a,b)$, then (a,b,c) is compatible if and only if c is in the shaded area.

Proof. We use an argument very similar to that in the proof of Corollary 10. We restrict our attention to the plane spanned by a, b and c and assume without loss of generality that $a=(0,-1)$, $b=(0,0)$ and that c is not to the left of the y-axis. By definition of compatibility, $c\notin B_x(d(x,a))$, where x is the element of $X(a,b)$ to the right of the y-axis. Under this condition, we have $\angle abc > \angle abc'$, where c' satisfies $d(c,b)=d(c',b)$ and is on the boundary of $B_x(d(x,a))$ above the x-axis. By definition of $X(a,b)$, we have $\cos\angle abx = \frac{\epsilon\sqrt{4-\epsilon^2}}{2}$, as illustrated in Figure 2b. Similar considerations show that $\cos\angle xbc' = \frac{\epsilon\sqrt{4-\epsilon^2}d(c',b)}{2d(a,b)}$. We have $d(c',b)=d(c,b)$ by assumption, and $\frac{\epsilon\sqrt{4-\epsilon^2}}{2}\le\frac{0.66\sqrt{4-0.66^2}}{2}<0.6231$. Since arccos is decreasing, this yields

$$
\begin{aligned}
\angle abc &> \angle abc' \\
&= \angle abx + \angle xbc' \\
&> \arccos(0.6231) + \arccos\left(0.6231\frac{d(c,b)}{d(a,b)}\right)
\end{aligned}
$$

and then the wanted inequality follows from $\arccos(0.6231) > 51.45°$.

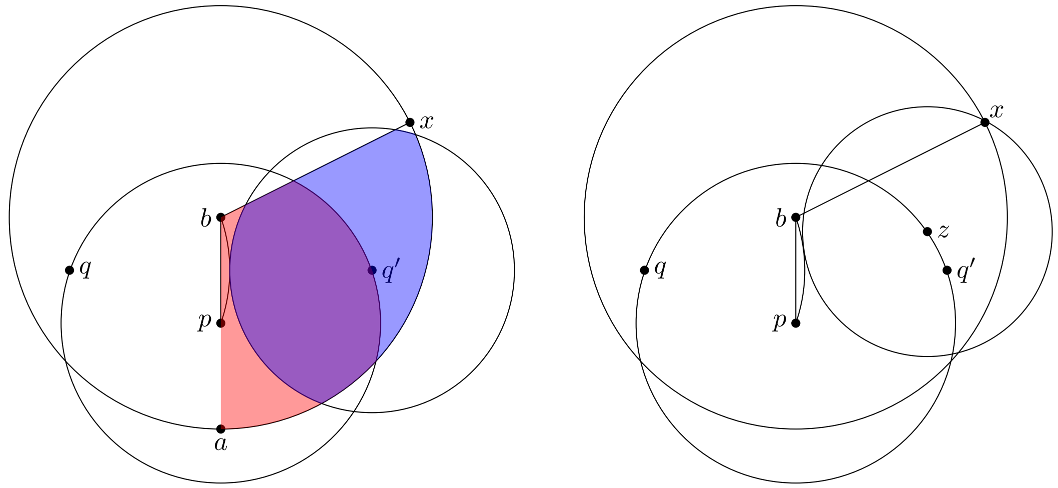

(a) The discs $B_p\left(\frac{d(p,b)}{\epsilon}\right)$ (red) and $B_{q'}(d(q,q')-\epsilon^{-1})$ (blue) cover the relevant area except a small part close to x.

(b) The distance from z to $B_q(\epsilon^{-1})$ is slightly smaller than $d(z,x)$.

Figure 4

(a,b,c) is compatible if and only (c,b,a) is, so the inequality holds also if we switch a and b. Thus, we can assume $d(a,b) \geq d(c,b)$. Under this assumption, the right-hand side is smallest when $d(c,b) = d(a,b)$. Thus,

$$\angle abc > 2\arccos(0.6231) > 102.9°. \quad \blacktriangleleft$$

3.3 The closest compatible neighbors are the correct neighbors

In the runtime analysis of our algorithms, we stated that checking if a triple (a,b,c) of points is compatible can be done in constant time. Since we only need to consider the geometry of three fixed points, this is clear; to be precise, by arguments similar to those in the proof of Lemma 13, what we need to check is if

$$\angle abc > \arccos\left(\frac{0.66\sqrt{4-0.66^2}}{2}\right) + \arccos\left(\frac{0.66\sqrt{4-0.66^2}d(c,b)}{2d(a,b)}\right)$$

and the same with a and c switching places.

Recall that our algorithms construct edges from $b \in \mathcal{S}$ to a and c, where a is the closest point in $\mathcal{S}$ to b, and c is the closest point in $\mathcal{S}$ to b such that (a,b,c) is compatible. (Compatible-crust is restricted to the Delaunay neighbors, which by Corollary 5 is not a problem.) Since b has exactly two adjacent vertices in $G_{\mathcal{C}}(\mathcal{S})$, it is sufficient to prove that a, b and c are consecutive. This is exactly the statement of Proposition 16 below, which therefore finishes the proof of Theorem 1.

▶ **Lemma 14.** *Let $\epsilon \leq 0.66$ and $a \to b$, let p be the midpoint on $[a,b]$, and let c be a point on $\mathcal{C} \setminus [a,b]$ with $d(b,c) \leq d(a,b)$. Then $\angle pbc > 117.3°$.*

Proof. Assume $\angle pbc \leq 117.3°$, and let us restrict ourselves to a plane containing p, b, c. Without loss of generality, we can assume that $p = (0,-1)$, $b = (0,0)$, and that c is not to the left of the y-axis. Let q be the point to the left of the y-axis such that $d(q,b) = d(q,p) =$

ϵ^{-1}, and let q' be the reflection of q across the y-axis. By Lemma 6 (i), $\mathrm{cl}(q') \in (p, b)$, and by Lemma 6 (ii) (choose $a = p$ in the lemma), $(p, b) \subset B_q(\epsilon^{-1})$. It follows that $c \notin B_{q'}(d(q, q') - \epsilon^{-1})$.

We have two remaining possibilities under the assumptions $\angle pbc \leq 117.3°$ and $d(b, c) \leq d(a, b)$:

(i) $c \in B_b(d(a, b)) \cap B_p\left(\frac{d(p,b)}{\epsilon}\right)$,

(ii) $c \in B_b(d(a, b)) \setminus \left(B_p\left(\frac{d(p,b)}{\epsilon}\right) \cup B_{q'}(d(q, q') - \epsilon^{-1})\right)$,

To show that (i) is impossible, first assume that c is below or on the line l through q and q'. If $\overline{cq}$ does not intersect $\overline{ab}$, let $I = \overline{cq}$. Otherwise, let $I = \overline{cq'}$. c is closer to any point on I than a is, so $a \notin \mathrm{cl}(I)$. Since no point on I is above l, $b \notin \mathrm{cl}(I)$, as p is always at least as close as b. But clearly, $\mathrm{cl}(c) = c$, and we have already observed that $\mathrm{cl}(q') \in (p, b)$, and $\mathrm{cl}(q) \in (p, b)$ holds for the same reason. Thus, $\mathrm{cl}(I)$ is disconnected, which contradicts the continuity of cl.

If instead c is above l, let $I = \overline{cq'}$ and use a similar argument with a and b exchanged.

Finally, we assume (ii), which is the case that requires the most care. Let $z = (1.244, -0.1351)$. Some calculation shows that $z \in B_p(\epsilon^{-1}) = B_p\left(\frac{d_p}{\epsilon}\right)$. Let $I = \overline{zq'}$. Since $I \subset B_p\left(\frac{d_p}{\epsilon}\right)$, I does not intersect the medial axis of C.

Let x be the intersection of the ray from b into the first quadrant with angle $\angle 117.3°$ with the boundary of $B_b(2)$. As Figure 4a illustrates, c must be in an area close to x, and x is the point in this area furthest away from z. Some more calculation shows that $d(z, x) < 1.18 < d(z, q) - \epsilon^{-1}$; see Figure 4b. This means that z is closer to c than to any point on $[p, b]$, since $[p, b] \subset B_q(\epsilon^{-1})$, as we have observed. Thus, $\mathrm{cl}(z) \notin [p, b]$. In addition, $\mathrm{cl}(q') \in (p, b)$ by Lemma 6 (i). But all points on I are closer to c than to both p and b (it is enough to check the endpoints of I), so $p, b \notin \mathrm{cl}(I)$. Thus, $\mathrm{cl}(I)$ is disconnected, which is impossible, as cl is continuous. ◀

▶ **Proposition 15.** *Let $\epsilon \leq 0.66$, and let a be a sample point and b a closest neighbor to a among the other sample points. Then a and b are consecutive.*

Proof. Suppose for a contradiction that x, a and y are consecutive and $b \notin \{x, y\}$. Let p be the midpoint of $[x, a]$ and q the midpoint of $[a, y]$. By Corollary 10, $\angle paq > 141°$, and by Lemma 14, both $\angle pab$ and $\angle qab$ are greater than $117.3°$, as $d(x, a), d(y, a) \geq d(a, b)$. The sum of these angles is greater than $360°$, which is impossible. ◀

▶ **Proposition 16.** *Let $\epsilon \leq 0.66$. Let b be a sample point, a a closest sample point to b, and c the closest sample point to b such that (a, b, c) is compatible. Then a, b and c are consecutive.*

In particular, there is a unique closest point c to a such that (a, b, c) is compatible.

The idea of the proof is as follows: We let a, b and c be as in the proposition, suppose there is a $c' \neq c$ such that a, b and c' are consecutive, let q be the midpoint of $[b, c']$, and carefully pick a point $p \in [a, b]$. We get lower bounds on $\angle qbc$ and $\angle pbq$ by Lemma 14 and Proposition 9 depending on the distances from q, c and p to b. This gives an upper bound on $\angle cbp$, which leads to a contradiction by an argument similar to the one in the proof of Lemma 14. However, the proof is complicated by the degrees of freedom we have in choosing the distances from the various points to b. We give the details in [8, Appendix A.4].

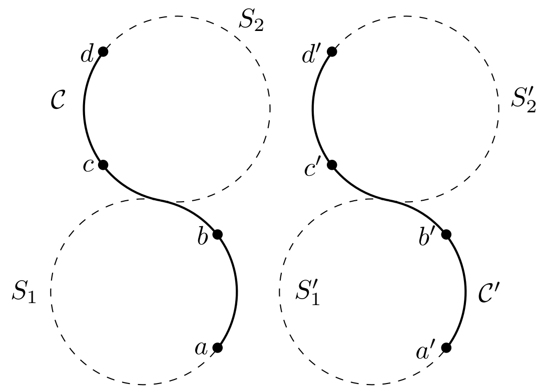

Figure 5 The first step in the construction of the curves of Theorem 17.

4 Counterexample to curve reconstruction for $\epsilon = 0.72$

In this section, we prove the following theorem, which says that correct curve reconstruction using 0.72-samples is not in general possible, even in $\mathbb{R}^2$. Moreover, one cannot determine whether the (union of) curve(s) has more than one connected component, and the reconstruction problem remains impossible also under the assumption that the sample is taken from a single connected curve.

▶ **Theorem 17.** *There is a finite set $\mathcal{S} \subset \mathbb{R}^2$ that is a 0.72-sample of $\mathcal{C}_1$, $\mathcal{C}_2$, $\mathcal{C}_3$ and $\mathcal{C}_4$, where $\mathcal{C}_1$ and $\mathcal{C}_2$ are connected closed curves and $\mathcal{C}_3$ and $\mathcal{C}_4$ are disconnected unions of closed curves, and $G_{\mathcal{C}_i}(\mathcal{S}) \neq G_{\mathcal{C}_j}(\mathcal{S})$ for all $i \neq j$.*

As we will construct subsets of the curves before we construct the complete curves, we extend the definition of ϵ-sampling to unions of closed curves in the obvious way.

Let $a = (0,-1)$, $b = (0,0)$, $c = (-1.008, 0.614)$, $d = (-1.008, 1.614)$. Let S_1 and S_2 be the two tangent circles with the same radius such that $a, b \in S_1$, $c, d \in S_2$ and the tangent point is the midpoint q between b and c. Let $\mathcal{C}$ be the union of the part of S_1 running from a to q through b and the part of S_2 running from q to d through c.

Next, let a', b', c', d' be the points, S_1', S_2' the circles and $\mathcal{C}'$ the curve we get by translating the whole construction horizontally to the right so that $d(b,c) = d(b,c')$; see Figure 5. Let T be the set of midpoints of $[a,b]$, $[b,c]$, $[c,d]$, $[a',b']$, $[b',c']$ and $[c',d']$.

▶ **Lemma 18.** *If $\{a,b,c,d,a',b',c',d'\}$ is not a 0.72-sample of $\mathcal{C} \cup \mathcal{C}'$, then there is a $t \in T$ such that $B_t\left(\frac{d_t}{0.72}\right)$ intersects the medial axis of $\mathcal{C} \cup \mathcal{C}'$.*

Proof. By definition, if $\{a,b,c,d,a',b',c',d'\}$ is not a 0.72-sample of $\mathcal{C} \cup \mathcal{C}'$, then there is a $p \in \mathcal{C} \cup \mathcal{C}'$ such that $B_p\left(\frac{d_p}{0.72}\right)$ intersects the medial axis of $\mathcal{C} \cup \mathcal{C}'$. Thus, it is enough to show that for every $p \in \mathcal{C} \cup \mathcal{C}' \setminus T$, there is a $t \in T$ such that $B_p\left(\frac{d_p}{0.72}\right) \subset B_t\left(\frac{d_t}{0.72}\right)$.

Let $p \in (x,y) \subset \mathcal{C}$ for some $x \to y$. We know that $d_p = d(p,x)$ or $d_p = d(p,y)$ by Lemma 4. Suppose $d_p = d(p,y)$ ($d_p = d(p,x)$ is similar), and pick $p' \in (p,y)$. Let $B = B_p\left(\frac{d_p}{0.72}\right)$ and $B' = B_{p'}\left(\frac{d_{p'}}{0.72}\right)$. If $B \nsubseteq B'$, there is a point on the ray from p through p' in $B' \setminus B$, which means that $\frac{d_p}{0.72} < d(p,p') + \frac{d_{p'}}{0.72}$, or equivalently

$$\frac{d_p - d_{p'}}{d(p,p')} < 0.72.$$

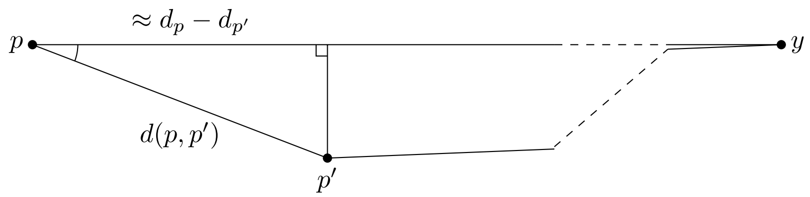

Figure 6 Assuming $d(p,p') \ll d(p,y)$, we have $\cos(\angle p'py) \approx \frac{d_p - d_{p'}}{d(p,p')}$. The dotted lines represent that we have collapsed a large part of the figure.

Observe that if we let p' approach p, then $\frac{d_p - d_{p'}}{d(p,p')}$ approaches $\cos\angle p'py$; see Figure 6. One can check that $\angle p'py < 40°$ for the possible p and p' in our example (by a large margin), while $\arccos(0.72) > 43°$. Thus,

$$\arccos(0.72) > \arccos\left(\frac{d_p - d_{p'}}{d(p,p')}\right)$$

for p' sufficiently close to p, so $0.72 < \frac{d_p - d_{p'}}{d(p,p')}$, a contradiction. This proves that as p moves along $\mathcal{C}$ or $\mathcal{C}'$ from a point in T towards a point in $\{a,b,c,d\}$, the disc $B_p\left(\frac{d_p}{0.72}\right)$ decreases (in the sense that later discs are contained in earlier discs), proving the lemma. ◀

▶ **Lemma 19.** *$\{a,b,c,d,a',b',c',d'\}$ is a 0.72-sample of $\mathcal{C} \cup \mathcal{C}'$.*

Proof. By Lemma 18, what we need to show is that for any $t \in T$, $B_t\left(\frac{d_t}{0.72}\right)$ does not intersect the medial axis.

To reduce the problem, observe that we have symmetry around the midpoint between b and c', as $\overrightarrow{bc} = -\overrightarrow{c'b'}$ and $\overrightarrow{cd} = -\overrightarrow{b'a'}$. Thus, we can restrict ourselves to the midpoints p, q and r of $[a,b]$, $[b,c]$ and $[c,d]$, respectively. For the rest of the proof, assume $t \in \{p,q,r\}$.

We extend cl to a set-valued map from $\mathbb{R}^2$ to $\mathcal{C} \cup \mathcal{C}'$ by letting $\mathrm{cl}(p)$ be the set of points in $\mathcal{C} \cup \mathcal{C}'$ that minimize the distance to p. Let m be a point such that $\mathrm{cl}(m)$ contains at least two points in $\mathcal{C} \cup \mathcal{C}'$, and let x and y be distinct points in $\mathrm{cl}(m)$. We will show that for $t \in \{p,q,r\}$, $m \notin B_t\left(\frac{d_t}{0.72}\right)$.

There are the following cases to consider:

- $x, y \in \mathcal{C}$,
- $x, y \in \mathcal{C}'$,
- $x \in \mathcal{C}$ and $y \in \mathcal{C}'$.

In the first case, m is on the medial axis of $\mathcal{C}$. This has two connected components: one is a curve starting at the center s_1 of S_1 and going leftwards and downwards from there, and the other is the mirror image through q of the first one. Because of symmetry, we only have to consider the first component. On this curve, s_1 minimizes the distance to p and q, and r is far away from the whole curve. One can check that the radius of S_1 is greater than 0.82, that $d_q = d(q,b) < 0.59$ and that $d_q > d_p$. Thus,

$$d(s_1,p) = d(s_1,q) > 0.82 > \frac{d_q}{0.72} > \frac{d_p}{0.72},$$

so we conclude that $B_t\left(\frac{d_t}{0.72}\right)$ does not intersect the medial axis of $\mathcal{C}$.

Next, we assume that $x, y \in \mathcal{C}'$. Then there is a point m' on the line segment $\overline{mt}$ such that $\mathrm{cl}(m')$ intersects both $\mathcal{C}$ and $\mathcal{C}'$, so m' is on the medial axis. If $m \in B_t\left(\frac{d_t}{0.72}\right)$, then $m' \in B_t\left(\frac{d_t}{0.72}\right)$, so we have reduced the second case to the third case.

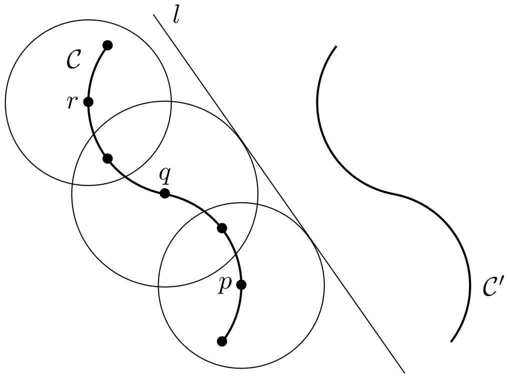

Figure 7 Illustration for the proof of Lemma 19.

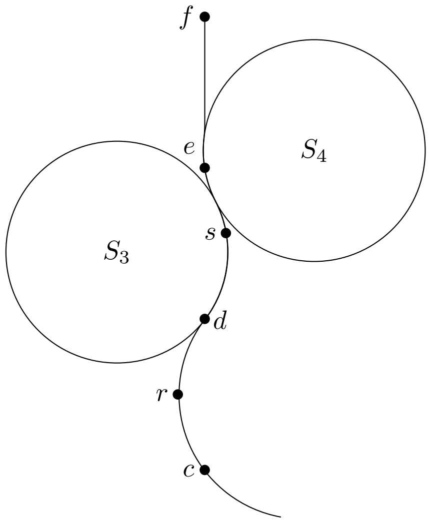

Figure 8 $\mathcal{C}$ is extended from d through e to f, where the tangent is vertical.

Lastly, assume that $x \in \mathcal{C}$ and $y \in \mathcal{C}'$; see Figure 7. Let l be the perpendicular bisector of s_1 and the center s_2' of S_2'. If $y \in S_2'$, then m is either on l or to the right of l (the latter can only happen if $x \in S_2$). By numerical calculation, one can check that $l \cap B_t\left(\frac{d_t}{0.72}\right) = \emptyset$, so in this case, $m \notin B_t\left(\frac{d_t}{0.72}\right)$. At the same time, if $y \in S_1'$, then $d(t, y) > \frac{2d_t}{0.72}$, so if $m \in B_t\left(\frac{d_t}{0.72}\right)$, then $y \notin S_1'$, as t is closer to m than S_1' is. ◀

Now we want to extend this construction. See Figure 8 for what follows. We add a point e such that d is the midpoint between c and e. Next, we put a circle S_3 with radius $\frac{d_r}{0.72}$ so that it is tangent to S_2 at d, and a circle S_4 with the same radius as S_3 tangent to S_3 such that e lies on S_4. If we extend $\mathcal{C}$ such that it contains $[d, e]$ along S_3 and S_4 in the obvious way, then $\{a, b, c, d, e\}$ is a 0.72-sample of $\mathcal{C}$. To see this, note that if s is the midpoint of $[d, e]$, then the difference in x-coordinate between s and d is less than that between r and d, so $d_s < d_r$. The closest points on the medial axis to s are the centers of S_3 and S_4, which have a distance of $\frac{d_r}{0.72} > \frac{d_s}{0.72}$ to s.

The tangent of $\mathcal{C}$ at d is much closer to being vertical than the tangent at e, and if we add another point f such that e is the midpoint between d and f, then we can extend $\mathcal{C}$ to f similarly to how we extended $\mathcal{C}$ from d to e in such a way that $\{a, b, c, d, e, f\}$ is a 0.72-sample, and such that the tangent of $\mathcal{C}$ at f is vertical.

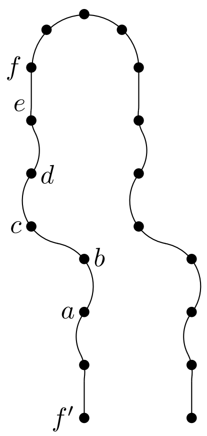

(a) $\mathcal{C}$ and $\mathcal{C}'$ are extended and tied together. At f and f', the tangents of the curve are vertical.

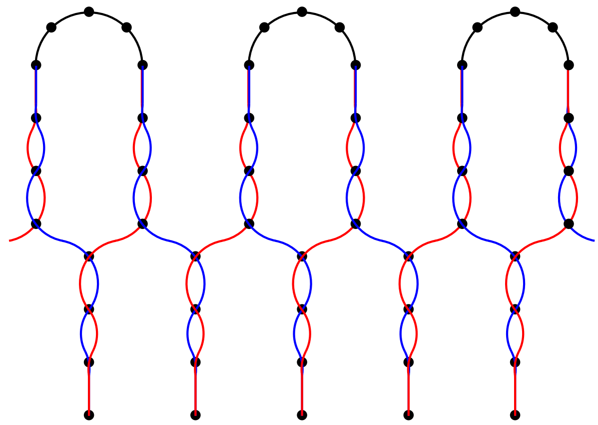

(b) Together with the black semicircles, the blue and red curves both give a valid reconstruction under the 0.72-sampling condition.

Figure 9

We can do the same below a, adding two points such that $\mathcal{C}$ can be extended downwards and the tangent of $\mathcal{C}$ at the lowest point is vertical. Now do the same for $\mathcal{C}'$, and add a sequence of points densely sampling a semicircle to connect $\mathcal{C}$ and $\mathcal{C}'$ as shown in Figure 9a. Again, the points shown make up a 0.72-sample of the curve. Next, we put many copies of this construction next to each other as shown in Figure 9b. Each copy is translated horizontally such that $d(b, c')$ is equal to the distance between b' in one copy and c in the copy on its right. If we ignore what happens to the far right or left, there are two ways to draw a set of curves with endpoints among the bottom points such that the set of points is a 0.72-sample of the union of curves.

We now take this long strip of points and curves and bend it slightly upwards such that they are contained in an annulus and the ends meet; see Figure 10. As the length of this strip goes to infinity, the distances from points on the curve to the closest sample point and the medial axis are distorted by a factor that approaches 1 when we bend it into the annulus. Our arguments for the the set of points being a 0.72-sample works equally well for an $\epsilon > 0.72$ sufficiently close to 0.72, so after turning the (sufficiently long) strip into an annulus, the point set stays a δ-sample for some $\delta > 0.72$.

Finally, we consider two such annuli with "the ends tied together", meaning that we draw curves between endpoints in the first annulus and endpoints in the second annulus, and sample the curves densely; see Figure 10. In each of the two annuli, we have two choices of how to draw the curve, as illustrated in Figure 9b, which gives four different choices. Exactly two of these choices result in a connected curve, and in all four cases, the set of points is a 0.72-sample of the curve or union of curves. Summing up, we get Theorem 17.

5 Counterexample to hypersurface reconstruction for $\epsilon = 0.72$

We have not defined what "correct reconstruction" means in higher dimensions. But assuming that preserving the number of connected components is required, we show that correct reconstruction of hypersurfaces in $\mathbb{R}^d$ using 0.72-samples is impossible for any $d \geq 2$.

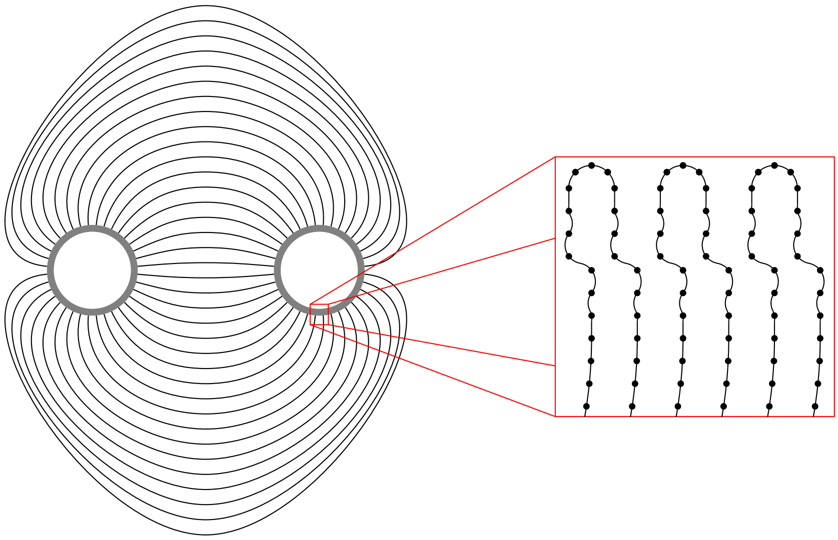

Figure 10 A reconstruction of the whole point set with the two annuli in grey. One can make sure that the curve is connected, and the point set is a 0.72-sample of it.

▶ **Theorem 20.** *For any $d \geq 2$, there is a finite point set $\mathcal{S} \subset \mathbb{R}^d$ that is a 0.72-sample of two manifolds $\mathcal{C}$ and $\mathcal{C}'$ without boundary of dimension $d-1$ with a different number of connected components.*

Proof. The case $d = 2$ follows immediately from Theorem 17. For any point $p = (x, y) \in (0, \infty) \times \mathbb{R}$, let p° be the circle centered at $(0, y)$ containing p. For any set $X \subset (0, \infty) \times \mathbb{R}$, let $X^\circ = \bigcup_{p \in X} p^\circ$. Let $\mathcal{C}_i$ be as in Theorem 17 for $1 \leq i \leq 4$, and let $\mathcal{S}_{\text{curve}}$ be the 0.72-sample as constructed in the previous section. Pick a constant R and translate $\mathcal{C}_i$ so that it is contained in $(R, \infty) \times \mathbb{R}$. Similarly to how we bent a strip into a large annulus earlier, by choosing R large, we can make sure that a sufficiently dense subset $\mathcal{S}$ of $\mathcal{S}^\circ_{\text{curve}}$ is a δ-sample of $\mathcal{C}_i$ for some $\delta > 0.72$. Choosing $i = 1$ and $i = 3$, the theorem for $d = 3$ follows. To get the theorem for larger d, one can iterate the construction we used to get from $d = 2$ to $d = 3$. ◀

6 Discussion

We have only considered unions of closed curves. An obvious question is if our work generalizes to open curves. We expect that this can be dealt with by a slight tweak of the algorithms when the endpoints are far apart: Instead of immediately connecting a point to its "correct" neighbors (i.e., its closest and closest "compatible" neighbors), one should add an edge between two points only when both points consider the other as a "correct" neighbor. However, we have not tried to turn this intuition into a precise statement.

Though this paper is mainly about curve reconstruction, we hope that it can also be a step towards improving the sampling conditions for surface reconstruction. Our arguments are valid for samples in any ambient dimension, and we expect many of our intermediate results to carry over to points on surfaces instead of curves. We consider generalizing our approach to surface reconstruction to be a promising direction of future research.

References

1 Ahmed Abdelkader, Chandrajit L. Bajaj, Mohamed S. Ebeida, Ahmed H. Mahmoud, Scott A. Mitchell, John D. Owens, and Ahmad A. Rushdi. Sampling Conditions for Conforming Voronoi Meshing by the VoroCrust Algorithm. In *34th International Symposium on Computational Geometry (SoCG 2018)*, pages 1:1–1:16, 2018.

2 Ernst Althaus and Kurt Mehlhorn. Traveling salesman-based curve reconstruction in polynomial time. *SIAM Journal on Computing*, 31(1):27–66, 2001.

3 Nina Amenta, Marshall Bern, and David Eppstein. The crust and the β-skeleton: Combinatorial curve reconstruction. *Graphical models and image processing*, 60(2):125–135, 1998.

4 Nina Amenta, Sunghee Choi, Tamal K. Dey, and Naveen Leekha. A simple algorithm for homeomorphic surface reconstruction. In *Proceedings of the sixteenth annual symposium on Computational geometry (SCG 2000)*, pages 213–222, 2000.

5 Dominique Attali. r-regular shape reconstruction from unorganized points. *Computational Geometry*, 10(4):239–247, 1998.

6 Matthew Berger, Andrea Tagliasacchi, Lee M. Seversky, Pierre Alliez, Gael Guennebaud, Joshua A. Levine, Andrei Sharf, and Claudio T. Silva. A survey of surface reconstruction from point clouds. *Computer Graphics Forum*, 36(1):301–329, 2017.

7 Fausto Bernardini and Chandrajit L. Bajaj. Sampling and reconstructing manifolds using alpha-shapes. In *Proceedings of the 9th Canadian Conference on Computational Geometry (CCCG 1997)*, pages 193–198, 1997.

8 Håvard Bakke Bjerkevik. Tighter bounds for reconstruction from ϵ-samples. *arXiv preprint v2*, 2022. arXiv:2112.03656.

9 Harry Blum. A transformation for extracting new descriptors of shape. In *Models for the Perception of Speech and Visual Form*, pages 362–380. MIT Press, Cambridge, 1967.

10 Frédéric Chazal and André Lieutier. Smooth manifold reconstruction from noisy and non-uniform approximation with guarantees. *Computational Geometry*, 40(2):156–170, 2008.

11 Mark de Berg, Otfried Cheong, Marc J. van Kreveld, and Mark H. Overmars. *Computational Geometry: Algorithms and Applications*. Springer, 3rd edition, 2008.

12 Luiz Henrique De Figueiredo and Jonas de Miranda Gomes. Computational morphology of curves. *The Visual Computer*, 11(2):105–112, 1994.

13 Tamal K. Dey. *Curve and Surface Reconstruction: Algorithms with Mathematical Analysis*, volume 23. Cambridge University Press, 2006.

14 Tamal K. Dey, Joachim Giesen, Edgar A. Ramos, and Bardia Sadri. Critical points of distance to an ε-sampling of a surface and flow-complex-based surface reconstruction. *International Journal of Computational Geometry & Applications*, 18(1–2):29–61, 2008.

15 Tamal K. Dey and Piyush Kumar. A simple provable algorithm for curve reconstruction. In *Proceedings of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 1999)*, volume 99, pages 893–894, 1999.

16 Tamal K. Dey, Kurt Mehlhorn, and Edgar A. Ramos. Curve reconstruction: Connecting dots with good reason. *Computational Geometry*, 15(4):229–244, 2000.

17 Tamal K. Dey and Rephael Wenger. Fast reconstruction of curves with sharp corners. *International Journal of Computational Geometry & Applications*, 12(05):353–400, 2002.

18 Christopher Gold. Crust and anti-crust: a one-step boundary and skeleton extraction algorithm. In *Proceedings of the fifteenth annual symposium on Computational geometry (SCG 1999)*, pages 189–196, 1999.

19 Jacob E. Goodman, Joseph O'Rourke, and Csaba D. Toth. *Handbook of Discrete and Computational Geometry*. CRC press, 3rd edition, 2017.

20 Tobias Lenz. How to sample and reconstruct curves with unusual features. In *EWCG: Proc. of the 22nd European Workshop on Computational Geometry*, pages 29–32. Citeseer, 2006.

21 Partha Niyogi, Stephen Smale, and Shmuel Weinberger. Finding the homology of submanifolds with high confidence from random samples. *Discrete & Computational Geometry*, 39(1-3):419–441, 2008.

22 Stefan Ohrhallinger, Scott A. Mitchell, and Michael Wimmer. Curve reconstruction with many fewer samples. *Computer Graphics Forum*, 35(5):167–176, 2016.

23 Stefan Ohrhallinger, Jiju Peethambaran, Amal D. Parakkat, Tamal K. Dey, and Ramanathan Muthuganapathy. 2d points curve reconstruction survey and benchmark. *Computer Graphics Forum*, 40(2):611–632, 2021.

24 Peer Stelldinger. Topologically correct surface reconstruction using alpha shapes and relations to ball-pivoting. In *19th International Conference on Pattern Recognition (ICPR 2008)*, pages 1–4, 2008.

25 Peer Stelldinger and Leonid Tcherniavski. Provably correct reconstruction of surfaces from sparse noisy samples. *Pattern Recognition*, 42(8):1650–1659, 2009.