Combinatorial curve reconstruction in Hilbert spaces:
A new sampling theory and an old result revisited

Daniel Freedman

Rensselaer Polytechnic Institute, Department of Computer Science, Troy, NY 12180, USA

Abstract

The goals of this paper are twofold. The first is to present a new sampling theory for curves, based on a new notion of local feature size. The properties of this new feature size are investigated, and are compared with the standard feature size definitions. The second goal is to revisit an existing algorithm for combinatorial curve reconstruction in spaces of arbitrary dimension, the Nearest Neighbour Crust of Dey and Kumar [Proc. ACM-SIAM Sympos. Discrete Algorithms, 1999, pp. 893–894], and to prove its validity under the new sampling conditions. Because the new sampling theory can imply less dense sampling, the new proof is, in some cases, stronger than that presented in [Proc. ACM-SIAM Sympos. Discrete Algorithms, 1999, pp. 893–894]. Also of interest are the techniques used to prove the theorem, as they are unlike those used used in the curve reconstruction literature to date. © 2002 Elsevier Science B.V. All rights reserved.

Keywords: Combinatorial curve reconstruction; Hilbert space

1. Introduction

The goals of this paper are twofold. The first is to present a new sampling theory for curves, based on a new notion of local feature size. The properties of this new feature size are investigated, and are compared with the standard feature size definitions employed in many papers on combinatorial curve and surface reconstruction, including [1]. The second goal is to revisit an existing algorithm for combinatorial curve reconstruction in spaces of arbitrary dimension, the Nearest Neighbour Crust of Dey and Kumar [8], and to prove its validity under the new sampling conditions. Because the new sampling theory can imply less dense sampling, the new proof is, in some cases, stronger than that presented in [8]. Also of interest in their own right are the techniques used to prove the theorems, as they are unlike those used in the curve reconstruction literature to date.

E-mail address: freedman@cs.rpi.edu (D. Freedman).

The problem of combinatorial curve reconstruction, in the sense of this paper, may be framed as follows. Given a curve embedded in a Hilbert Space of arbitrary dimension, and N sample points from this curve, is it possible to construct a new curve based only on knowledge of the samples which is topologically equivalent to the original curve, and is sufficiently close to it? Hidden in the problem statement are the three principal difficulties that inhere in the problem. First is the combinatorial nature of the problem: the points are *unorganized*, that is, they may come in any order to the user. This fact precludes any straightforward reasoning about the adjacency relationships between samples. It may be argued that some topological and/or geometrical constraints on the curve to be reconstructed may be used to help alleviate this problem. This, in fact, is not the case; for the second obstacle associated with the problem of curve reconstruction is that no prior topological or geometrical information is available. Rather, both the topology and geometry of the underlying curve are assumed to be unknown; all available information about the curve is contained in the samples. The final hurdle is that the algorithm must work for embedding spaces of arbitrary dimension. The algorithm of Dey and Kumar, unlike several of its predecessors (see Section 4), is designed to work not only for curves embedded in the plane, but for curves living in any embedding space which is a Hilbert Space. Indeed, the embedding space need not even be finite-dimensional; it may, for example, be the space of functions $L_2(\mathbb{R})$.

Of course, such a problem may not be solved without some additional assumptions on the curve samples. In particular, a common assumption pertains to the fineness with which the curve is sampled. A standard approach involves defining a local feature size, $\phi(\cdot)$, which is a function mapping the curve to real numbers greater than 0. Intuitively, at any given point, this feature size should indicate roughly the size of the curve's features at that point. Thus, in relatively flat parts of the curve, ϕ will be quite high, while in areas with greater curvature, or which pass close to other parts of the curve, we expect that ϕ will be much smaller. Given such a ϕ, then for $0 < r < 1$, X is said to be an r-sampling of a curve C if (a) $X \subset C$ and (b) $\forall c \in C, \exists x \in X$ such that $\|c - x\| < r\phi(c)$. This notion of sampling is plausible, since areas which are relatively featureless can be sampled sparsely, while areas with more features must be sampled relatively more densely. An algorithm for curve reconstruction is usually established given such a sampling condition, namely, assuming that the samples form an r-sample of the curve C, and r is less than some critical sampling parameter r^*.

As a result, the particular definition of the feature size ϕ plays an important role in both the statement of the theorem, i.e., under what conditions the algorithm will yield a provably correct reconstruction, as well as the shape of the proof. In this paper, a new definition of the local feature size ϕ will be given. This definition is geometrically plausible, and will be discussed at length in Section 3. It will also be compared with the standard definition, used, for example, in [3,8]. In Section 4, we briefly review the Nearest Neighbour Crust of Dey and Kumar. In Section 5, we present the proof of the algorithm's validity, using the new definition of feature size; in consequence, the techniques used to prove the result differ greatly from those in [8]. It may be the case that such methods are of use in proving related results, such as combinatorial surface and manifold reconstruction. Finally, in Section 6 we show some results of applying the algorithm, and discuss aspects of complexity.

2. Previous work

The problem of combinatorial curve reconstruction has been discussed for some time, with early works focusing on various heuristics [7,11,17]. Algorithms which offer provably correct reconstructions were

first studied in the case in which sampling was uniform [5,6]. Amenta et al. [3] provided the first provably correct reconstructions, the crust and the β-skeleton, in the case of non-uniform sampling, using the more general sampling condition mentioned in Section 1. (Note that the prior work of Melkemi [15] contains similar results.) Other algorithms use a similar sampling condition, but present modified algorithms: that of Dey et al. [9] is designed to provide greater self-consistency, while that of Gold and Snoeyink [13] leads to a more conservative reconstruction. The paper of Dey and Kumar [8], which is discussed in this paper, generalizes curve reconstruction in a different way, by allowing for the embedding dimension to be general; all of the other algorithms work only for curves embedded in the plane.

All of the algorithms mentioned thus far are able to reconstruct smooth curves. More recent efforts have focused on the reconstruction of curves with cusps. Giesen [12] presented the first work which addresses the problem of cusps; he showed that a global approach, which sought to reconstruct the curve by finding the TSP tour amongst the samples, would lead to a provably correct reconstruction. Althaus and Melhorn [1] made this problem tractable by showing that the TSP could be found in polynomial time in this case. The deficiency of these approaches is that they assume that the curves are simply connected; using a new algorithm which does not have provable properties, Dey and Wenger [10] achieve good experimental results for multiply connected non-smooth curves. Finally, there is the somewhat unrelated "minimal interpolant" algorithm of Petitjean and Boyer [18], which assumes that the underlying curve itself must be a simplicial complex, thereby leading to a form of self-consistency.

Finally, there exists the related work on combinatorial surface reconstruction. In addition to some of the work in the field of graphics, such as [14], the computational geometry literature provides several examples, such as the algorithms of Amenta and Bern [2], and the somewhat simpler algorithm by Amenta et al. [4].

3. The sampling theory

Before explaining the new definition of local feature size, we must introduce some notation. The embedding space is Z, a real Hilbert Space. The inner product on Z is denoted $\langle \cdot, \cdot \rangle : Z \times Z \to \mathbb{R}$, and the norm is $\|z\| = \langle z, z \rangle^{1/2}$. The results to be presented, if appropriately modified in minor ways, would also apply to complex Hilbert Spaces; however, the extra notation tends to obfuscate the results. A curve C is a $\mathcal{C}^1$ 1-manifold embedded in Z. As in [3], it is assumed that C has neither endpoints nor self-intersections. The samples of the curve are denoted $X \subset C$, with $|X| = N$. The curve has m connected components; the ith has arc-length L_i, and a unit-speed parameterization given by $f_i : \mathbb{R} \to Z$, where f_i is $\mathcal{C}^1$ and periodic with period L_i. Thus, C may be written

$$C = \bigcup_{i=1}^{m} \{ z \in Z : z = f_i(u), u \in [0, L_i] \}.$$

The tangent at a point c on the curve is denoted $t(c)$, which is equal to $f_i'(u)$ where $c = f_i(u)$; note that $\|t(c)\| = 1$. A normal to the curve at c is denoted $n(c)$, and is also generally assumed to have unit norm.

A ray between the points z_1 and z_2 is denoted $\rho(z_1, z_2)$ and is the vector $z_2 - z_1$. Let $\alpha(z_1, z_2)$ be the angle between two vectors z_1 and z_2. Given $p, \delta \in Z$, interpret p as a point and δ as a direction; then a *beam* is defined by

$$beam(p, \delta) = \{ z \in Z : |\alpha(\rho(p, z), \delta)| \leqslant \beta \}$$

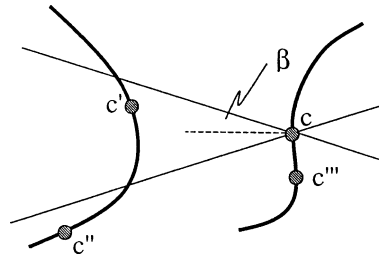where β is referred to as the *beam-width*.

Fig. 1. Illumination. Portions of a curve are illustrated in bold; the normal to the curve is dashed; and the boundary of the beam is shown in thin solid lines. The point c' is illuminated by the point c; the points c'' and c''' are not.

We may now turn to the task of defining the feature size; this definition rests on the concepts of *illumination* and its close relative, *visibility*. A point $c_2 \in C$ is said to be *illuminated* by $c_1 \in C$ if there exists a *beam*$(c_1, n(c_1)) \ni c_2$. c_1 is said to be *visible* to c_2 if c_1 illuminates c_2 or c_2 illuminates c_1. Note that visibility is symmetric: c_1 is visible to $c_2 \Leftrightarrow c_2$ is visible to c_1. For an illustration of the concept of illumination, see Fig. 1. Now, let Θ be the visibility set,

$$\Theta(c) = \{c' \in C : c' \text{ is visible to } c\}.$$

Then we may define the *visible distance* at c by

$$\theta(c) = \min_{c' \in \Theta(c)} \|c - c'\|.$$

Now, defining

$$\Phi(c) = \left\{ c' \in C : \|c' - c\| \leqslant \theta(c) \right\}$$

leads to our definition of *local feature size*,

$$\phi(c) = \min_{c' \in \Phi(c)} \theta(c').$$

Before investigating any more formal properties of $\phi(c)$, let us spend some time developing the intuition for why this quantity respects features. One can imagine, at a particular point c on the curve, shining a beam of light which illuminates other parts of the curve; see Fig. 1. The beam, which subtends an angle 2β, can be thought of as tethered to the curve at c using a hinge which can rotate about the curve; thus, the beam can shine in any direction which is normal to the curve at the point from which it emanates. The points which are illuminated by the beams are said to be illuminated by c. Looking at things from the opposite vantage point, we may wonder about which points illuminate c. In combination, all of these points are said to be visible to c: either c can shine its beam on one of these points, or one of these points can shine its beam on c.

Intuitively, the proximity of visible points to c determine how featured this part of the curve is. Fig. 2 demonstrates how the notion of visible distance, $\theta(\cdot)$, captures features. In Fig. 2(a), a feature near the point c is shown. This feature has a source which is "local to c", i.e., it is caused by the curvature of the curve near c. This kind of feature will be captured by the notion of visible distance, as c' is illuminated by c, and thus we have that $\theta(c) \leqslant \|c - c'\|$. Figs. 2(b) and 2(c) show examples of features whose sources are not "local" to c, but rather are "global", i.e., are due to completely different parts of the curve. In Fig. 2(b), the source of the feature is a different part of the same connected component, whereas in Fig. 2(c), the source is a different connected component. Once again, the relevant feature is illuminated by c, and thus
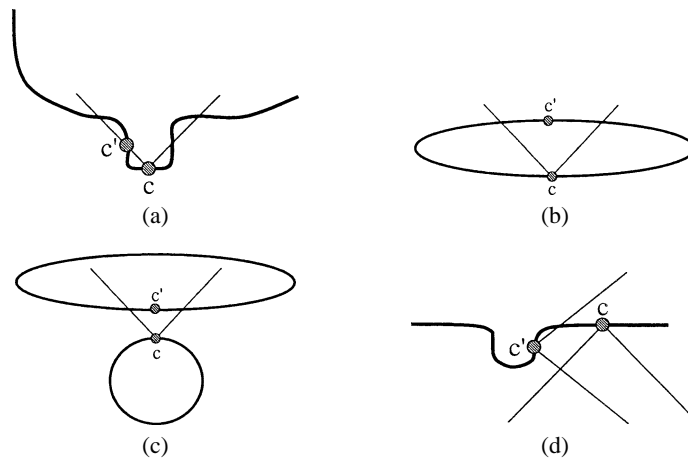
Fig. 2. Examples of visible points, illustrating the concept of visible distance. See accompanying discussion in the text.

the concept of visible distance captures the proximity of this feature. Thus far, all examples have shown the feature being illuminated by the point c; Fig. 2(d) shows an example where this doesn't happen. If the notion of visible distance at c were based solely on the points which are illuminated by c, then this feature would not be captured. However, note that c is illuminated by the point c' on the feature; thus, c' is visible to c, and the feature is in fact captured by the visible distance. This is the reason that the notion of visibility is defined symmetrically.

Thus far, we have discussed the intuitive meaning of the visible distance $\theta(\cdot)$. The question now arises: why not define the local feature size as $\theta(\cdot)$? The answer is that, largely for technical reasons, it is necessary for the local feature size to be Lipschitz, i.e., it must satisfy

$$\left| \phi(c_1) - \phi(c_2) \right| \leqslant K \| c_1 - c_2 \|$$

for all points c_1 and c_2 on the curve, and some $K > 0$. While θ does not possess this property, ϕ does. $\phi(c)$ is found taking the smallest value of $\theta(\cdot)$ in a small neighbourhood of the curve around c; interestingly, the size of the neighbourhood itself depends on $\theta(c)$.

Let us compare some of the properties of the new feature size $\phi(\cdot)$ introduced in this paper, with the more common notion of feature size, denoted $\xi(\cdot)$, used in [8] (and by many other papers as well). Note that both this paper, as well as that of Dey and Kumar [8], establish homeomorphic reconstruction assuming $\frac{1}{3}$-sampling. It is thus quite important to know the relative sizes of ϕ and ξ. For example, if it were the case that $\xi(c) \geqslant \phi(c)$ at all points on all curves, then the results established by Dey and Kumar [8] would be strictly better than those established here, as they would require less dense sampling to achieve the appropriate reconstruction. What we shall show now is that this is not the case; rather, there are cases in which $\xi(c) > \phi(c)$, and there are also cases in which $\phi(c) > \xi(c)$. Let us examine each of these cases in turn.

Focus on the curve $C = \{z \in \mathbb{R}^2 : \|z\| = 1\}$, i.e., the unit circle in the plane. Clearly, the medial axis is just the single point $(0, 0)$. Thus, $\xi(c) = 1$ for all points c on the curve. Computing $\phi(c)$ is somewhat more complex; let us begin by examining the set of points which are illuminated by a given c. For
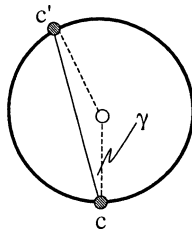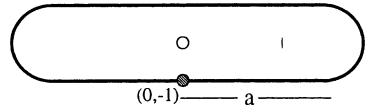
Fig. 3. For a circle, we have that $\xi(c) > \phi(c)$.

Fig. 4. For an oval, we have that $\phi(0, -1) > \xi(0, -1)$. If $a \gg 1$, then $\phi(c) = 2\xi(c)$ for most points on the oval.

concreteness, fix $c = (0, -1)$. Referring to Fig. 3, we can see that for the visible point c', we must have that

$$\|c - c'\| = \big\|(-1, 0) - (-\sin 2\gamma, \cos 2\gamma)\big\| = \sqrt{2(1 + \cos 2\gamma)}.$$

This function is monotonically decreasing in γ, and we know that any illuminated point c' must have $|\gamma| \leqslant \beta$, where β is the beam-width. Thus,

$$\min_{|\gamma| \leqslant \beta} \|c - c'\| = \sqrt{2(1 + \cos 2\beta)}.$$

Due to the circle's symmetry, it is easy to show that the only points which illuminate c are those which are illuminated by c. Thus,

$$\theta(c) = \sqrt{2(1 + \cos 2\beta)}.$$

Once again, due to the circle's symmetry, we must have that $\theta(c)$ is the same for all c, which leads finally to

$$\phi(c) = \sqrt{2(1 + \cos 2\beta)}.$$

In order to compare $\phi(c)$ with $\xi(c)$, we need to know something more about the beam-width β. It will turn out that appropriate reconstruction is only provably guaranteed with $\beta \geqslant 74°$. In this case, we have that

$$\phi(c) \leqslant 0.55$$

and therefore, $\xi(c) > \phi(c)$ for all points c on the circle.

The situation is quite different for the curve illustrated in Fig. 4. This figure shows an oval, i.e., a curve defined as

$$C = \big\{z \in \mathbb{R}^2 \colon |z_1| \leqslant a, z_2 = 1\big\} \cup \big\{z \in \mathbb{R}^2 \colon |z_1| \leqslant a, z_2 = -1\big\}$$
$$\cup \big\{z \in \mathbb{R}^2 \colon \|z - (a, 0)\| \leqslant 1, z_1 \geqslant a\big\} \cup \big\{z \in \mathbb{R}^2 \colon \|z - (-a, 0)\| \leqslant 1, z_1 \leqslant -a\big\}.$$

In other words, the oval is defined as two line segments and two semi-circles. a is a parameter which gives the lengths of the line segments. The medial axis is quite easy to compute, and is just the set $\{z \in \mathbb{R}^2 \colon |z_1| \leqslant a, z_2 = 0\}$. Thus, as in the case of the circle, $\xi(c) = 1$ for all points c on the oval. Now, focus on the point $\hat{c} = (0, -1)$. If $a \geqslant 2$, then the closest point that $\hat{c}$ illuminates if the point $(0, 1)$. A brief inspection also shows that there is no point which is closer than $(0, 1)$ which illuminates $\hat{c}$. Thus,

$$\theta(\hat{c}) = \big\|(0, -1) - (0, 1)\big\| = 2.$$

Furthermore, if $a \geqslant 4$, then all points c within a distance of $\theta(\hat{c}) = 2$ away from $\hat{c}$ (namely, the points $c = (u, -1)$ with $|u| \leqslant 2$) will also have $\theta(c) = 2$, by the argument used above. Therefore, if $a \geqslant 4$, we have that $\phi(\hat{c}) = 2$. In other words, $\phi(\hat{c}) = 2\xi(\hat{c})$; the new notion of feature size is twice as big (at $\hat{c}$) as the standard definition, and therefore would lead to less dense sampling. Indeed, if a is quite large, i.e., the oval is very long and thin, then most points c on the oval (all except those near the semi-circular caps) will satisfy $\phi(c) = 2\xi(c)$.

To summarize, then, there are points on curves for which $\xi(c) \geqslant \phi(c)$, and there are points on curves for which $\phi(c) \geqslant \xi(c)$. Since this paper shows that the NN-crust algorithm is valid for $\frac{1}{3}$-sampling using the feature size $\phi(c)$, we have in some cases improved upon the result of Dey and Kumar [8], which shows that the algorithm is valid for $\frac{1}{3}$-sampling using the feature size $\xi(c)$. Combining the results of [8] and those of this paper, we have the following. Let $\Lambda(c) = \max\{\phi(c), \xi(c)\}$; then the NN-crust produces a homeomorphic reconstruction of a curve C using $\frac{1}{3}$-sampling with the feature size given by $\Lambda(c)$. This is a strict improvement over [8].

4. The algorithm

The algorithm, which Dey and Kumar refer to as the "Nearest Neighbour Crust", is very simple, and is described in Fig. 5.

The curve is reconstructed as a simplicial complex, i.e., as a polygonal approximation to the underlying curve. The algorithm therefore tries to find the two sample points which are adjacent to a given sample. The idea is that one of these points is always the point which is closest to the sample in question. The second adjacent point is, in some sense, the sample which is "on the other side" of the original. Thus, if a is the direction from the sample to its closest point, the other adjacent point will be an x' which satisfies $\langle a, x' - x \rangle < 0$; it must lie in the opposite halfspace. Indeed, it is taken to be the closest such point.

Note that this formulation relies on the fact that the curve C is $\mathcal{C}^1$. If the curve were merely $\mathcal{C}^0$, then at cusps, the two adjacent points could very well lie in the same direction.

The algorithm proposed is very simple, both to understand and to implement. However, it is quite another matter to establish the validity of the algorithm, namely to demonstrate that the reconstruction which it provides is both homeomorphic to the original curve, as well as sufficiently close to it. The next

$\hat{C} = \text{RECONSTRUCT}(X)$
  $\hat{C} = \emptyset$
  for each $x \in X$
    $\hat{x} = \text{argmin}_{x' \in X - \{x\}} \|x' - x\|$
    $a = \hat{x} - x$
    $X_o = \{x' \in X: \langle a, x' - x \rangle < 0\}$
    $\tilde{x} = \text{argmin}_{x' \in X_o} \|x' - x\|$
    $\hat{\ell} = \{z \in Z: z = (1 - \lambda)x + \lambda\hat{x}, \lambda \in [0, 1]\}$
    $\tilde{\ell} = \{z \in Z: z = (1 - \lambda)x + \lambda\tilde{x}, \lambda \in [0, 1]\}$
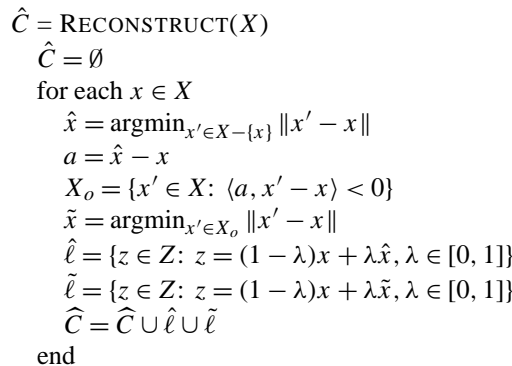 $\widehat{C} = \widehat{C} \cup \hat{\ell} \cup \tilde{\ell}$
  end

Fig. 5. The reconstruction algorithm.

section addresses these issues in a formal context, and shows under what conditions RECONSTRUCT leads to a proper reconstruction of the original curve.

5. Proof of the algorithm

For concreteness, let us focus on one particular sample point $x_1 \in X$. Let x_2 and x_3 be adjacent to x_1. Furthermore, without loss of generality, let $\|x_2 - x_1\| \leqslant \|x_3 - x_1\|$. Assume that x_1, x_2, and x_3 lie on the ith connected component of C. Let the parameters corresponding to x_k be u_k; assume, without loss of generality, that $u_3 < u_1 < u_2$. (The orientation of the parameterization can always be flipped to accommodate this result.) Also, let $d_j(u) = \|f_j(u) - x_1\|$.

In what is to follow, the main goal will be to demonstrate that $\hat{x} = x_2$ and $\tilde{x} = x_3$ if sampling is sufficiently dense; i.e., that the two points chosen by RECONSTRUCT are adjacent to x_1. Lemmas 1 and 2 give a particular partition of the points on the ith component of the curve. This partition has the three useful properties outlined in Lemma 2, properties which will be utilized throughout the proof.

Lemma 1. $\exists u^+, u^-$ *such that* (a) $u_1 < u^+ \leqslant u^- < u_1 + L_i$ *and* (b) $d_i'(u^+) = d_i'(u^-) = 0$.

Proof. The derivative d_i' may be calculated as

$$d_i'(u) = \frac{\langle f_i'(u), f_i(u) - x_1 \rangle}{\|f_i(u) - x_1\|}$$

which exists everywhere but at $u = u_1$. Since f_i is $\mathcal{C}^1$, d_i' is everywhere continuous, except at $u = u_1$. Now, it can be shown that

$$\lim_{u \to u_1 + L_i:\ u < u_1 + L_i} d_i'(u) = -1, \qquad \lim_{u \to u_1:\ u > u_1} d_i'(u) = 1.$$

Since d_i' is continuous, by the intermediate value theorem there exists at least one $w \in (u_1, u_1 + L_i)$ such that $d_i'(w) = 0$. Let u^+ be the smallest such w, and let u^- be the largest such; then by definition, we have that $d_i'(u^+) = d_i'(u^-) = 0$, and that $u_1 < u^+ \leqslant u^- < u_1 + L_i$. (Note that $u^+ = u^-$ is a possibility, if there is exactly one such w.) $\square$

Lemma 2. *Using the definition of u^+ and u^- from Lemma 1, the set $(u_1, u_1 + L_i)$ can be partitioned into $\{\mathfrak{L}^+, \mathfrak{N}, \mathfrak{L}^-\}$, where $\mathfrak{L}^+ = (u_1, u^+)$, $\mathfrak{N} = [u^+, u^-]$, and $\mathfrak{L}^- = (u^-, u^- + L_i)$. This partition has the following properties*:

(1) $u \in \mathfrak{L}^+ \Rightarrow d_i'(u) > 0$;
(2) $u \in \mathfrak{L}^- \Rightarrow d_i'(u) < 0$;
(3) $u^* = \operatorname{argmin}_{u \in \mathfrak{N}} d_i(u)$ *is well-defined, and satisfies* $d_i'(u^*) = 0$.

Proof. Properties (1) and (2) follow directly from the definitions of u^+ and u^-, see Lemma 1. To prove property (3), note that $\mathfrak{N}$ is a compact set, and that $d_i(u)$ is continuous with respect to u (since C is $\mathcal{C}^1$); this implies that u^* is well-defined, see for example [16]. If $u^* \in (u^+, u^-)$, then $d_i'(u^*) = 0$; a global minimum which is in the interior of a set must also be a local minimum. If $u^* = u^+$, we have

a boundary solution, for which u^* is no longer necessarily a local minimum; however, by the definition of u^+, $d_i'(u^+) = 0$ (see Lemma 1). A similar argument applies in the case of u^-, thereby completing the proof. □

Lemmas 3 and 4 establish the fact that all local minima of the $d_j(\cdot)$ functions must be visible to x_1.

Lemma 3. *Let β be the beam-width, and let*

$$\psi(c, c') = \frac{|\langle c' - c, t(c) \rangle|}{\|c' - c\|}.$$

If $\gamma = \sin(\beta)$, then

$$\psi(c, c') \geqslant \gamma \quad and \quad \psi(c', c) \geqslant \gamma \quad \Leftrightarrow \quad c' \notin \Theta(c).$$

Proof. If $c' \notin \Theta(c)$, then c does not illuminate c', and c' does not illuminate c. Look at the first case. For any normal $n(c)$, it must be that $|\alpha(c' - c, n(c))| > \beta$. In particular, this must be true for the normal which lies in the plane defined by placing the tails of the vectors $t(c)$ and $c' - c$ at the point c. This implies that $|\alpha(c' - c, t(c))| \leqslant \frac{\pi}{2} - \beta$. Finally,

$$\cos\left(\alpha(c' - c, t(c))\right) = \frac{\psi(c, c')}{\|t(c)\|} = \psi(c, c')$$

since C has a unit-speed parameterization. This in turn yields that $\psi(c, c') \geqslant \cos(\pi/2 - \beta) = \sin(\beta)$. The argument for the second case is exactly analogous. □

Lemma 4. *If u^* is a local minimum of $d_j(\cdot)$ for any j, and $c^* = f_j(u^*)$ then $c^* \in \Theta(x_1)$.*

Proof. If u^* is a local minimum of $d_j(\cdot)$, we must have that $d_j'(u^*) = 0$, i.e.,

$$\frac{\langle f_j'(u^*), f_j(u^*) - x_1 \rangle}{\|f_j(u^*) - x_1\|} = 0.$$

Noting that $t(c^*) = f_j'(u^*)$, and using the converse of Lemma 3 completes the proof. □

Lemma 5 shows that distance from x_1 to any point on the jth component on the curve, $j \neq i$, must be greater than the visible distance $\theta(x_1)$. Furthermore, the distance from x_1 to any point which lies in the part of the ith component of the curve given by $f_i(\mathfrak{N})$, must also be greater than the visible distance.

Lemma 5. *Let $\mathfrak{M} = f_i(\mathfrak{L}^+) \cup f_i(\mathfrak{L}^-) \cup \{x_1\}$. Then if $c \notin \mathfrak{M}$, $\|c - x_1\| \geqslant \theta(x_1)$.*

Proof. First, examine the case when c is on the ith component of C. In this case, $c \notin \mathfrak{M} \Rightarrow c \in f_i(\mathfrak{N})$. From Lemma 2, we know that $\exists u^* \in \mathfrak{N}$ such that $d_i(u^*) \leqslant d_i(u) \ \forall u \in \mathfrak{N}$ and $d_i'(u^*) = 0$. Thus, $\|c - x_1\| \geqslant d_i(u^*)$. From Lemma 4,

$$d_i'(u^*) = 0 \ \Rightarrow \ f_i(u^*) \in \Theta(x_1) \ \Rightarrow \ d_i(u^*) \geqslant \theta(x_1).$$

Thus, $\|c - x_1\| \geqslant \theta(x_1)$.

Now consider the case when c is on the jth component of C for $j \neq i$. Let u^* be the global minimum of $d_j(\cdot)$; then $\|c - x_1\| \geqslant d_j(u^*)$. Because d_j is both $\mathcal{C}^1$ and periodic, the global minimum must also be a local minimum. Thus, $f_j(u^*) \in \Theta(x_1)$. Using the same argument as above, we may then conclude that $\|c - x_1\| \geqslant \theta(x_1)$. $\quad \square$

The essential result given in Lemma 5 can now be used to show that the closest sample $\hat{x}$ to x_1 must lie in the ith component, either in $f_i(\mathfrak{L}^+)$ or in $f_i(\mathfrak{L}^-)$. Further, it can be shown that $\hat{x} = x_2$. These results are given in Lemmas 6–8, and Theorem 1.

Lemma 6. $\exists x_{12} \in C$ such that $\|x_1 - x_{12}\| = \|x_2 - x_{12}\|$. x_{12} lies on the same component as x_1 and x_2, between them.

Proof. Let
$$\delta(u) = \big\| f_i(u) - f_i(u_1) \big\| - \big\| f_i(u) - f_i(u_2) \big\|.$$
Since f_i is continuous, so is δ. Let $\eta = \|f_i(u_2) - f_i(u_1)\|$; then $\delta(u_2) = \eta$, while $\delta(u_1) = -\eta$. By the intermediate value theorem $\exists u_{12} \in (u_1, u_2)$ such that $\delta(u_{12}) = 0$; then if $x_{12} = f_i(u_{12})$, $\|x_1 - x_{12}\| = \|x_2 - x_{12}\|$. Note that x_{12} is *not* a sample, i.e., $x_{12} \notin X$. The notation x_{12} is merely used as a mnemonic for the reader to remember x_{12}'s property. $\quad \square$

Lemma 7. If X is an r-sampling of C for any $r \leqslant 1$, $\|x_1 - x_2\| < 2r\phi(x_{12})$.

Proof. To establish this result, use Lemmas 1, 2, 4, and 5, substituting the functions $d_j(u) = \| f_j(u) - x_{12}\|$ (rather than $\|f_j(u) - x_1\|$). In this case, all points on the curve which do not belong to $\mathfrak{M}(x_{12})$ satisfy $\|c - x_{12}\| \geqslant \theta(x_{12}) \geqslant \phi(x_{12})$. Since X is an r-sampling of C, there must exist $x \in X$ such that $\|x_{12} - x\| < r\phi(x_{12})$; by the previous statement, such an x cannot be outside $\mathfrak{M}(x_{12})$ as long as $r < 1$. That is, there must exist an $x \in \mathfrak{M}(x_{12})$ which is closer to x_{12} than are any of the samples which lie outside of $\mathfrak{M}(x_{12})$. Now, from Lemma 2, $d_i'(u) > 0 \ \forall u \in \mathfrak{L}^+(x_{12})$, and $d_i'(u) < 0 \ \forall u \in \mathfrak{L}^-(x_{12})$. As a result, any x in $\mathfrak{M}(x_{12})$ which minimizes d_i must also minimize $|u - u_{12}|$. There are two possible candidates, x_1 and x_2; and $\|x_1 - x_{12}\| = \|x_2 - x_{12}\|$. Since both x_1 and x_2 are closest to x_{12}, by the r-sampling condition we must have that $\|x_i - x_{12}\| < r\phi(x_{12})$ for $i = 1, 2$. Finally, by the triangle inequality, $\|x_1 - x_2\| \leqslant \|x_1 - x_{12}\| + \|x_{12} - x_2\| < 2r\phi(x_{12})$. $\quad \square$

Lemma 8. If X is an r-sample of C with $r \leqslant 1$, then $\phi(x_{12}) \leqslant \theta(x_1)$.

Proof. $\phi(c) = \min_{c' \in \Phi(c)} \theta(c')$, where
$$\Phi(c) = \big\{c' \in C: \|c' - c\| \leqslant \theta(c)\big\}.$$
Thus, $\phi(x_{12}) \leqslant \theta(c)$ for all c such that $\|c - x_{12}\| \leqslant \theta(x_{12})$. But in proving Lemma 7, we established that $\|x_1 - x_{12}\| < r\phi(x_{12}) \leqslant \phi(x_{12}) \leqslant \theta(x_{12})$. Thus, $\phi(x_{12}) \leqslant \theta(x_1)$. $\quad \square$

Theorem 1. If X is an r-sample of C for $r \leqslant \frac{1}{2}$, then $\|x_2 - x_1\| \leqslant \|x - x_1\| \ \forall x \in X - \{x_1, x_2\}$.

Proof. There are two cases to consider: $x \in \mathfrak{M}$ and $x \notin \mathfrak{M}$. For the former case, we have noted from Lemma 2, $d_i'(u) > 0 \ \forall u \in \mathfrak{L}^+(x_1)$, and $d_i'(u) < 0 \ \forall u \in \mathfrak{L}^-(x_1)$. As a result, any x in $\mathfrak{M}$ which

minimizes d_i must also minimize $|u - u_1|$. There are two possible candidates, x_2 and x_3; and by assumption $\|x_1 - x_2\| \leqslant \|x_1 - x_3\|$, so $x_2 = \mathrm{argmin}_{x \in \mathfrak{M}} \|x - x_1\|$. Now, consider the case $x \notin \mathfrak{M}$: such an x was shown to satisfy $\|x - x_1\| \geqslant \theta(x_1)$ (see Lemma 5). However, from Lemma 7, we know that $\|x_1 - x_2\| < 2r\phi(x_{12}) \leqslant 2r\theta(x_1)$, where the latter follows from Lemma 8. However, since $r \leqslant \frac{1}{2}$, then $2r \leqslant 1$, and thus $\|x_1 - x_2\| < \theta(x_1)$. This completes the proof. $\quad\square$

Lemma 9 shows that x_3 is closer to x_1 than all points except possibly those which lie inside $f_i(\mathfrak{L}^+)$. This allows future results to concentrate on the region $f_i(\mathfrak{L}^+)$.

Lemma 9. *If $r \leqslant \frac{1}{2}$, x_3 is closer to x_1 than all $x \notin \mathfrak{M}$ and all $x \in \mathfrak{L}^-$.*

Proof. Using most of the previous lemmata, but focusing on x_3 instead of x_2 (and substituting x_{13} for x_{12}) shows that $\|x_3 - x_1\| < \theta(x_1)$. Thus, x_3 is closer to x_1 than is any sample $x \notin \mathfrak{M}$. It has also been noted in Lemma 2 that within $\mathfrak{M}$, $d_i'(u) < 0 \ \forall u \in \mathfrak{L}^-$; since u_3 is closer to u_1 than is the u-parameter of any other sample within $\mathfrak{L}^-$, we must have that x_3 is closer to x_1 than these points as well. $\quad\square$

Lemmas 10 and 11 show that x_3 does indeed lie inside X_o, that is, $\langle x_3 - x_1, x_2 - x_1 \rangle < 0$. Thus, x_3 is a candidate to be $\tilde{x}$.

Lemma 10. *Let $a = \frac{x_2 - x_1}{\|x_2 - x_1\|}$. For X an r-sample of C with $r \leqslant \frac{1}{2}$, we have that $\|a - t(x_1)\| \leqslant \sqrt{2(1 - \gamma)}$.*

Proof. We have shown in Theorem 1 that if $r \leqslant \frac{1}{2}$, then $\|x_2 - x_1\| < \theta(x_1)$. Thus, $x_2 \notin \Theta(x_1)$; using Lemma 3, this implies that

$$\frac{|\langle x_2 - x_1, t(x_1) \rangle|}{\|x_2 - x_1\|} = |\langle a, t(x_1) \rangle| \geqslant \gamma.$$

Thus, either $\langle a, t(x_1) \rangle \geqslant \gamma$ or $\langle a, t(x_1) \rangle \leqslant -\gamma$. To show that $\langle a, t(x_1) \rangle \geqslant \gamma$, assume that $\langle a, t(x_1) \rangle \leqslant -\gamma$ and establish a contradiction. Let

$$Q(u) = \frac{|\langle f_i(u) - x_1, t(x_1) \rangle|}{\|f_i(u) - x_1\|}.$$

Then $Q(u_2) = \langle a, t(x_1) \rangle \leqslant -\gamma$. Now, it is easy to show that

$$\lim_{u \to u_1: \ u > u_1} Q(u) = 1.$$

Also, since $f_i(\cdot)$ is continuous, then $Q(\cdot)$ must be continuous. Thus, by the intermediate value theorem, $\exists \hat{u} \in (u_1, u_2)$ such that $Q(\hat{u}) = 0$. Using the converse of Lemma 3, $Q(\hat{u}) = 0 \Rightarrow \hat{x} = f_i(\hat{u}) \in \Theta(x_1)$. Thus, $d_i(\hat{u}) \geqslant \theta(x_1)$. However, since $\hat{u} \in (u_1, u_2)$, any $u \in [\hat{u}, u_2]$ satisfies $d_i'(u) > 0$, since all such $u \in \mathfrak{L}^+$. Thus, $d_i(u_2) > d_i(\hat{u}) \geqslant \theta(x_1)$. This is a contradiction, since we have already established that for $r \leqslant \frac{1}{2}$, $d_i(u_2) < \theta(x_1)$. Thus, we must have that $\langle a, t(x_1) \rangle \geqslant \gamma$.

Finally, $\|a - t(x_1)\|^2 = \|a\|^2 + \|t(x_1)\|^2 - 2\langle a, t(x_1) \rangle$. Using the fact that $\|a\| = \|t(x_1)\| = 1$, and the above result completes the proof. $\quad\square$

Lemma 11. *If the beam-width $\beta \geqslant 74°$, then $\langle x_3 - x_1, x_2 - x_1 \rangle < 0$.*

Proof. Let $b = \frac{x_3 - x_1}{\|x_3 - x_1\|}$. Since $\|x_3 - x_1\| < \theta(x_1)$, we may use the reasoning in Lemma 10, applied to x_3, to show that $\|b + t(x_1)\| \leqslant \sqrt{2(1 - \gamma)}$. (Note that we must compare b to $-t(x_1)$, which accounts for the plus sign above.) Now, using the triangle inequality note that $\|t(x_1) - (-t(x_1))\| \leqslant \|t(x_1) - b\| + \|b + t(x_1)\|$, so that $\|b - t(x_1)\| \geqslant 2 - \|b + t(x_1)\|$ (since $\|t(x_1)\| = 1$) or $\|b - t(x_1)\| \geqslant 2 - \sqrt{2(1 - \gamma)}$. Now, again use the triangle inequality to find that $\|b - t\| \leqslant \|b - a\| + \|a - t\|$, so that

$$\|a - b\| \geqslant \|b - t\| - \|a - t\| \geqslant 2 - 2\sqrt{2(1 - \gamma)}.$$

Finally, note that since $\|a\| = \|b\| = 1$,

$$\langle a, b \rangle = 1 - \frac{1}{2}\|a - b\|^2 \leqslant -5 + 4\gamma + 4\sqrt{2(1 - \gamma)}.$$

Since $\beta \geqslant 74° \Rightarrow \gamma = \sin(\beta) > 0.96$, which gives that $\langle a, b \rangle < -0.028$. $\quad\square$

Theorem 2 establishes the fact that $\tilde{x} = x_3$. When combined with Theorem 1, this establishes the fact that the algorithm picks out both points correctly, i.e., both points which are adjacent to x_1. Finally, Theorem 3 formalizes this result, to show that the curve generated by the RECONSTRUCT algorithm is indeed homeomorphic to the original curve.

Theorem 2. *Let* $a = x_2 - x_1$ *and* $X_o = \{x \in X: \langle a, x - x_1 \rangle < 0\}$. *Then if* $r \leqslant \frac{1}{2}$, *and the beamwidth* $\beta \geqslant 49°$, *then* $x_3 = \operatorname{argmin}_{x \in X_o} \|x - x_1\|$.

Proof. From Lemma 9, the only samples we need to consider are those $x \in \mathfrak{L}^+$. It may be that for such x, we have $\langle a, x - x_1 \rangle \geqslant 0$, so that none of these x belong to X_o; in this case, the proof is completed. However, suppose that this is not the case, and there exists an $x \in \mathfrak{L}^+$ such that $\langle a, x - x_1 \rangle < 0$.

Note that $\langle a, x_2 - x_1 \rangle = \|x_2 - x_1\| > 0$, and the function $\delta(u) = \langle a, f_i(u) - x_1 \rangle$ is continuous; thus, before $\delta(u)$ becomes negative, it must pass through 0. Let us therefore examine an $x^* = f_i(u^*)$ such that $x^* \in \mathfrak{L}^+$, $u^* > u_2$, and $\langle a, x^* - x_1 \rangle = 0$. Define $h = (x^* - x_1)/\|x^* - x_1\|$. Then

$$\langle t, h \rangle = 1 - \frac{1}{2}\|t - h\|^2,$$

since $\|t\| = \|h\| = 1$. Also,

$$\|t - h\| \geqslant \|a - h\| - \|a - t\|.$$

Further, $\|a - h\|^2 = 2 - 2\langle h, a \rangle = 2$ (recall that $\langle h, a \rangle = 0$ is the definition of x^*). Also, from Lemma 10, $\|a - t(x_1)\| \leqslant \sqrt{2(1 - \gamma)}$. Thus,

$$\|t - h\| \geqslant \sqrt{2} - \sqrt{2(1 - \gamma)},$$

which gives that $\langle t, h \rangle \leqslant 2\sqrt{1 - \gamma} - (1 - \gamma)$. Using $\beta \geqslant 49° \Rightarrow \gamma = \sin(\beta) > 0.754$, so that $0 < \langle t, h \rangle < 0.746$.

Finally, you can see that $|\langle t, h \rangle| < \gamma$, so that from Lemma 3, $x^* \in \Theta(x_1)$. Of course, for any $u \geqslant u^*$, we will have $\|x - x_1\| \geqslant \theta(x_1)$, since $d_i'(u) > 0$ for $u \in \mathfrak{L}^+$. Thus, we are done since $\|x_3 - x_1\| \leqslant \theta(x_1)$. $\quad\square$

Theorem 3. *If* X *is an* r-*sample of* C *for* $r \leqslant 1/3$, *and the beamwidth* $\beta \geqslant 74°$, *then* $\widehat{C} =$ RECONSTRUCT(X) *is homeomorphic to* C.

Proof. From Theorem 1, we have that $\hat{x} = x_2$, where $\hat{x} = \operatorname{argmin}_{x \in X - \{x_1\}} \|x - x_1\|$; that is, the closest point to x_1 is also adjacent to it. From Theorem 2, $\tilde{x} = x_3$, where $\tilde{x} = \operatorname{argmin}_{x \in X_o} \|x - x_1\|$; that is, the point which is "on the other side" of x_1 is also adjacent to x_1. Thus, we have established that only correct simplices are added to $\hat{C}$: simplices in which the edges points are adjacent in the underlying curve C. Now, let $g : C \to \hat{C}$ be defined as follows. For any two adjacent points $x_a = f_j(u_a)$ and $x_b = f_j(u_b)$, let g be defined on the portion of the curve $f_j([u_a, u_b])$ as follows:

$$g\big(f_j((1 - \lambda)u_a + \lambda u_b)\big) = (1 - \lambda)f_j(u_a) + (1 - \lambda)f_j(u_b)$$

for each $\lambda \in [0, 1]$. It is clear that g is a bijection. Continuity of g and g^{-1} are properties which follow directly from the definition of g. Thus, C is homeomorphic to $\hat{C}$.

It is interesting to note once again that the condition of $r \leqslant 1/3$ is precisely the same condition as in [8], albeit with a completely different definition of local feature size. $\quad\square$

6. Complexity and experiments

It remains to discuss issues of complexity. Looking at the structure of the pseudocode in Fig. 5, it is clear that the algorithm is $O(N^2)$ for $N = |X|$. Although this is not quite as good as the $O(N \log N)$ complexity of [3], it is nice that the complexity is independent of the embedding dimension. (In fact, if the sampling condition is the same as that of [3], the algorithm can be shown to run in $O(N \log N)$ in the case of curves embedded in the plane; see [8].)

The algorithm has been applied in four cases. Unfortunately, it is very difficult to show images of curves that are embedded in four or higher dimensions; thus, we will have to content ourselves with experimental results concerning curves embedded in the plane and in space (although, of course, the algorithm is valid for curves embedded in arbitrary Hilbert Spaces). Figs. 6 and 7 show the samples and corresponding reconstructions for curves embedded in the plane; Figs. 8 and 9 show results for space. Note the varying topologies and geometries that are presented.

Note that in Figs. 6 and 8, the curves are not closed, as is supposed in the proof of the algorithm's correctness. However, the algorithm is specially modified in such circumstances: the distances
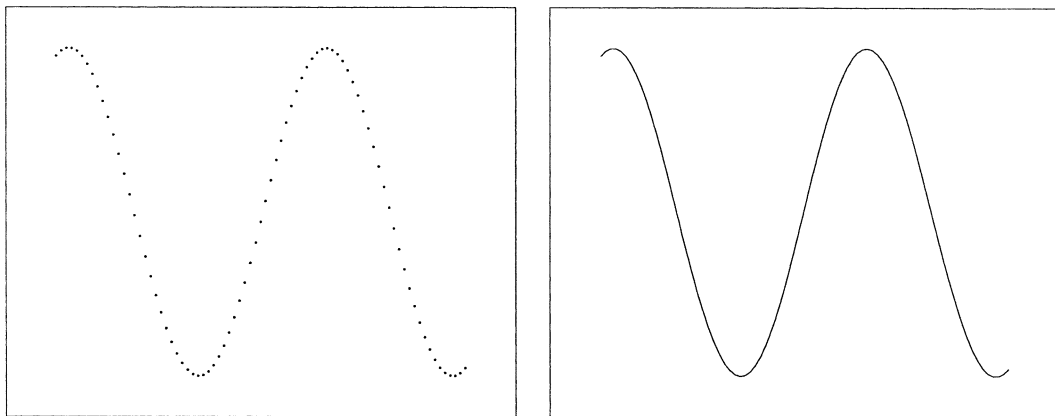

Fig. 6. Curve reconstruction using the RECONSTRUCT algorithm, for a simple curve embedded in two dimensions. Note that in this case, the curve has endpoints, so a minor modification to the algorithm is employed; see text.
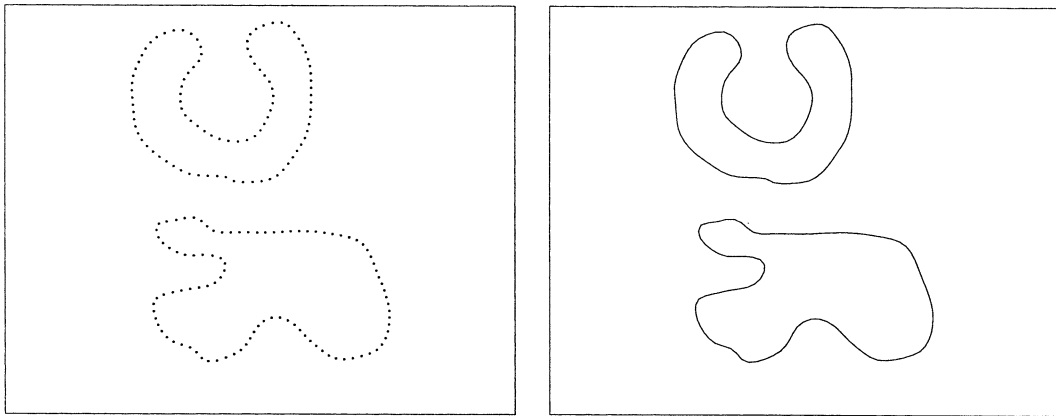
Fig. 7. Curve reconstruction for a curve embedded in two dimensions. Note that this curve has more complex geometry and topology.
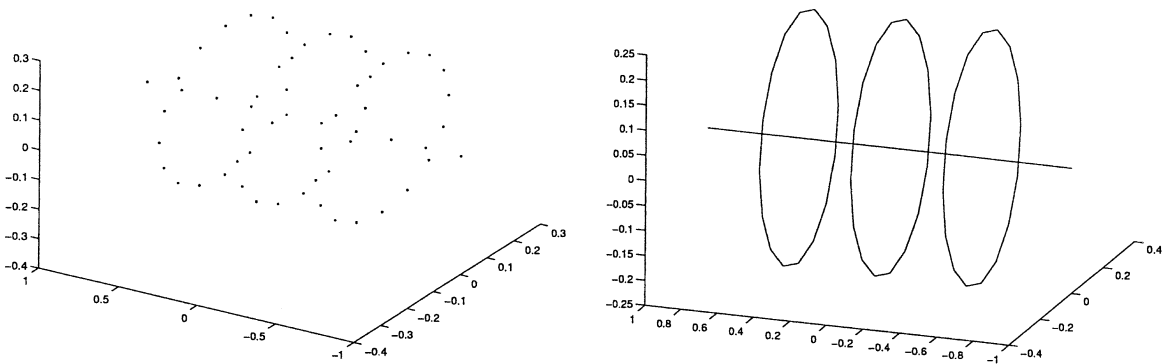

Fig. 8. Curve reconstruction for a curve embedded in three dimensions. Note that the curve has endpoints, so a minor modification to the algorithm is employed; see text.
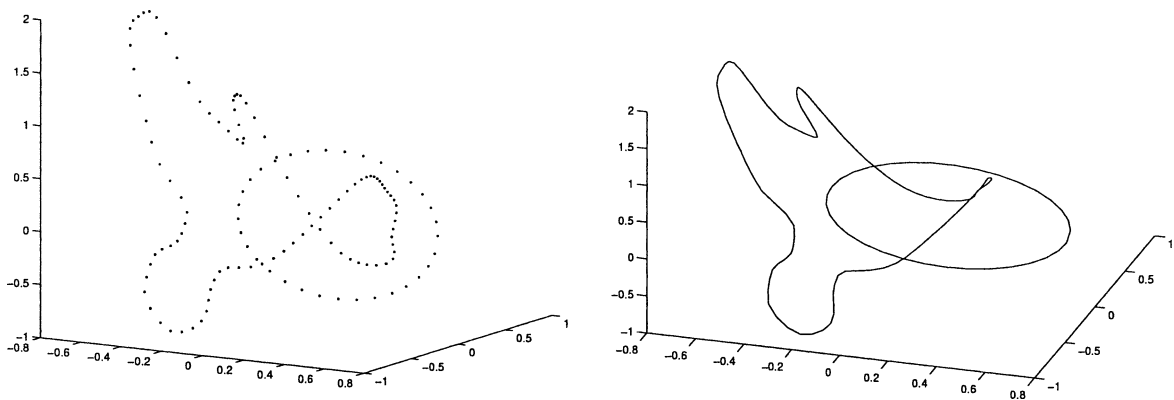

Fig. 9. Curve reconstruction for a complicated curve embedded in three dimensions.

$\hat{d} = \|x - \hat{x}\|$ and $\tilde{d} = \|x - \tilde{x}\|$ are compared. If $\tilde{d} > \alpha\hat{d}$ for some fixed constant $\alpha > 1$, then the segment between x and $\tilde{x}$ is not included in the reconstruction. A value of $\alpha = 3$ worked well in most cases.

References

[1] E. Althaus, K. Mehlhorn, Polynomial time tsp-based curve reconstruction, in: Proc. 11th ACM-SIAM Sympos. Discrete Algorithms, 2000, pp. 686–695.

[2] N. Amenta, M. Bern, Surface reconstruction by Voronoi filtering, Discrete Comput. Geom. 22 (1999) 481–504.

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

[4] N. Amenta, S. Choi, T. Dey, N. Leekha, A simple algorithm for homeomorphic surface reconstruction, in: Proc. 16th ACM Sympos. Comput. Geom., 2000, pp. 213–222.

[5] D. Attali, r-regular shape reconstruction from unorganized points, in: Proc. 13th ACM Sympos. Comput. Geom., 1997, pp. 248–253.

[6] F. Bernardini, C. Bajaj, Sampling and reconstructing manifolds using α-shapes, in: Proc. 9th Canadian Conf. Comput. Geom., 1997, pp. 193–198.

[7] L. de Figueiredo, J. de Miranda Gomes, Computational morphology of curves, Visual Comput. 11 (1995) 105–112.

[8] T. Dey, P. Kumar, A simple provable algorithm for curve reconstruction, in: Proc. ACM-SIAM Sympos. Discrete Algorithms, 1999, pp. 893–894.

[9] T. Dey, K. Mehlhorn, E. Ramos, Curve reconstruction: Connecting dots with good reason, Computational Geometry 15 (2000) 893–894.

[10] T. Dey, R. Wenger, Reconstructing curves with sharp corners, Computational Geometry 19 (2) (2001) 89–99.

[11] H. Edelsbrunner, D.G. Kirkpatrick, R. Seidel, On the shape of a set of points in the plane, IEEE Trans. Inform. Theory 29 (4) (1983) 551–559.

[12] J. Giesen, Curve reconstruction, the tsp, and mengers theorem on length, in: Proc. 15th ACM Sympos. Comput. Geom., 1999, pp. 207–216.

[13] C. Gold, J. Snoeyink, A one-step crust and skeleton extraction algorithm, Algorithmica 30 (2) (2001) 144–163.

[14] H. Hoppe, T. DeRose, T. Duchamp, H. Jin, J. McDonald, W. Stuetzle, Piecewise smooth surface reconstruction, in: Proc. SIGGRAPH '94, 1994, pp. 19–26.

[15] M. Melkemi, $\mathcal{A}$-shapes and their derivatives, in: Proc. 13th ACM Sympos. Comput. Geom., 1997, pp. 367–369.

[16] J. Munkres, Topology: A First Course, Prentice-Hall, Englewood Cliffs, 1975.

[17] J. O'Rourke, H. Booth, R. Washington, Connect-the-dots: A new heuristic, Comput. Vision Graph. Image Process. 39 (1987) 258–266.

[18] S. Petitjean, E. Boyer, Regular and non-regular point sets: Properties and reconstruction, Computational Geometry 19 (2) (2001) 101–126.