How to Sample and Reconstruct Curves With Unusual Features

Tobias Lenz
Institut für Informatik, Freie Universität Berlin, Takustr. 9, 14195 Berlin, Germany
tlenz@mi.fu-berlin.de

Abstract

This work generalizes the ideas in the Nearest-Neighbor-Crust algorithm by Dey and Kumar. It allows to reconstruct smooth closed curves from ε-samples with $\varepsilon \leq 0.48$ which is a big improvement compared to the original bound. Further generalization leads to a new algorithm which reconstructs arbitrary curves (open, closed, smooth, with corners, with intersections) in any dimension. The algorithm works well in practice but lacks a nice sampling condition comparable to the well-known ε-sampling condition. This shortcoming is posed as an open problem.

1 Introduction

In curve reconstruction, one cannot or does not want to transfer a complex shaped curve. Instead, a finite *sample* of the original curve is generated which must become reconstructed at the destination. A reconstruction is just a polygonal line connecting the points in the correct order. The approximation quality measured as distance to the original can be arbitrarily bad depending on the sample. It is obvious, that not every point set leads to the correct reconstruction, therefore the sample must have special properties called the *sampling condition*.

2 Picking Samples from a Curve

Definition 1 *A* sample *S is a finite subset of a curve Σ. The elements of the sample are points in $\mathbb{R}^d$ and are called* sample points.

The ε-sampling condition basically states that the sample should be dense if the curvature is high or parts of the curve are close together while it might be loose in straight regions of the curve. Formally it is defined as follows.

Definition 2 *The* medial axis *of a smooth curve is the set M containing the center points of all empty circles which touch the curve in more than one point.*

The local feature size *of a point $p \in \Sigma$ is defined as $lfs(p) = \min_{m \in M} \|p - m\|$.*

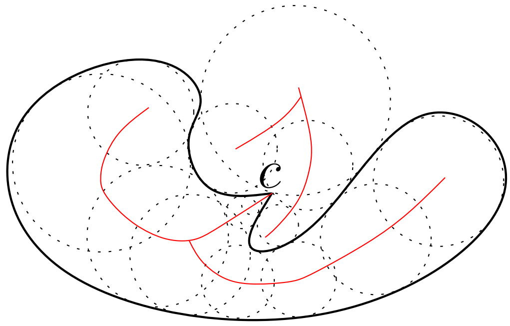

Figure 1: A non-smooth curve (bold, black) with some medial balls (dotted) and the medial axis (red) intersecting the curve at the corner point c.

Definition 3 *A sample S is an* ε-sample *of the smooth curve Σ if and only if $\forall p \in \Sigma : \exists s \in S : \|p - s\| \leq \varepsilon \cdot lfs(p)$.*

These definitions apply well to smooth curves but an ε-sample is impossible for curves with corners or intersections because the medial axis goes through these points and hence the local feature size becomes zero. Therefore the sampling density should be infinite. See figure 1 for an illustration. A possible solution is to exclude a small neighborhood of these points from the ε-sampling condition and define a new one for these regions. This is used for corners in [2].

3 The NN-Crust Algorithm

The basis for this work came from Dey and Kumar and their *Nearest-Neighbor-Crust* algorithm [1]. They start with a point set and connect each point p to its nearest neighbor u and its nearest half-neighbor h. Consider the line perpendicular to $\overline{pu}$ through p which splits the plane. The half-neighbor h of p is now the point closest to p which does not lie in the half plane containing u.

This procedure results in a set of edges forming the reconstruction. If the underlying point set is at least a $\frac{1}{3}$-sample of a smooth closed curve, the NN-Crust algorithm guarantees a correct reconstruction. Dey and Kumar showed that the edges in the reconstruction are a subset of the edges of a Delaunay triangulation of the point set. This allows simple implementations

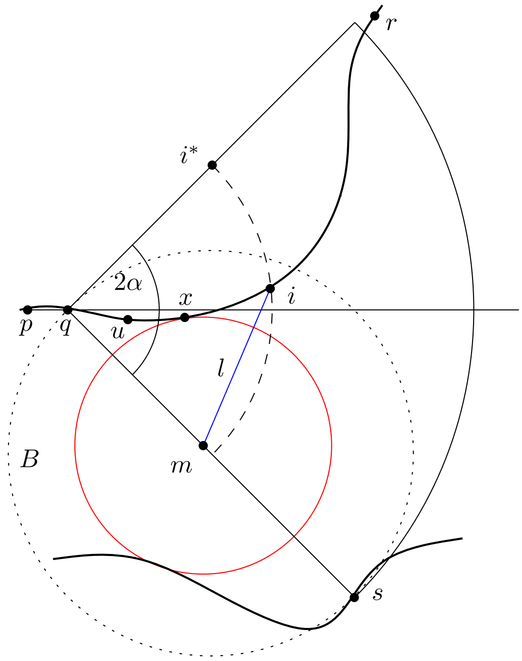

Figure 2: Computing the worst case distance between a point on the curve and the medial axis.

with $\mathcal{O}(n \log n)$ running time in 2d and extensions to higher dimensions.

4 Generalization of the NN-Crust Algorithm

4.1 Flexible Opening Angles

Changing the way of describing the half-neighbor slightly allows a simple generalization. A similar definition to the one given in section 3 is the following: The half-neighbor of a point p with nearest neighbor u in a point set S is the point $h \in S \setminus \{p\}$ which minimizes the distance between p and h and fulfills $\angle \overline{up}, \overline{ph} > \frac{\pi}{2}$. Now the angle appears as a parameter and instead of $\frac{\pi}{2}$ one can choose an arbitrary angle. These points are no longer called "half-neighbors" but *α-neighbors* instead, where α specifies the maximally allowed turning angle going from u to p to h.

The remainder of this section is used to show that for α-neighbors with $0 < \alpha \leq \frac{\pi}{2}$, the resulting edges are still a correct reconstruction. The arguments refine and extend ideas from the original paper which leads to a much better bound for ε. The proofs are skipped due to space constraints.

Lemma 1 *The angle spanned by three adjacent samples in an ε-sample with $\varepsilon < 1$ is at least $\pi - 4 \arcsin \frac{\varepsilon}{2}$. This angle corresponds to $\pi - \alpha$, which gives $\varepsilon \leq 2 \sin \frac{\alpha}{4}$ if α is fixed and less than $\frac{\pi}{2}$.*

Figure 2 shows how to bound ε from below using the following arguments.

Lemma 2 *Given three adjacent sample points p, q, r in an ε-sample with $\varepsilon < 1$, the curve segment between q and r runs completely inside the cone at q aligned with (p, q) with opening angle $2\alpha = 8 \arcsin \frac{\varepsilon}{2}$.*

Lemma 3 *Consider a $\frac{1}{2}$-sample. Given a chain of three adjacent samples p, q, r and a sample s not adjacent to q. If p and r do not lie inside the ball B with diameter $\overline{qs}$ there will be a medial axis point on the segment $\overline{qs}$ with distance at most $\frac{3}{4}\|s - q\|$ from q.*

Lemma 4 *Given an ε-sample of a smooth closed curve with $\varepsilon \leq 0.48$ and a correct edge (p, q). The edge (q, s) is correct if and only if s is the closest point to q inside the cone with apex q aligned to (p, q) with opening angle $2 \cdot 0.97$.*

Altogether this guarantees correct results for a whole family of algorithms with different angles $0 < \alpha \leq \frac{\pi}{2}$ and for $\varepsilon \leq 0.48$ in the extreme case which is a big improvement compared to the NN-Crust [1]. The NN-Crust algorithm is a special case of this approach for $\alpha = \frac{\pi}{2}$. It follows a direct increase of the ε bound for the original algorithm from $\frac{1}{3}$ to 0.4 due to a more careful analysis.

4.2 Non-Circular Neighborhoods

Picking the nearest neighbor essentially is like growing a circle around a point until another point touches the circle's boundary. Instead of growing circles one can use other shapes, here called *probes*, for example with the intention to prefer straight segments over bends or left turns over right turns. This requires an alignment of the shape because it is no longer rotationally symmetric. Since every vertex in a reconstruction of a smooth closed curve has degree two, one only has to find a single *seed edge* and grow the reconstruction from that, aligning the shape at the tip of an already reconstructed edge. A good choice for the seed edge is the overall shortest edge which is obviously a correct edge for ε-samples.

Compared to the global approach of triangulating a point set and then selecting a subset of the triangulation edges, tracing is a more local concept. Extending the reconstructed graph at its loose ends with minimum weight components is also a very natural procedure for greedy algorithms. This idea is in line with famous algorithms like Dijkstra's shortest path algorithm or Prim's algorithm to construct a minimum spanning tree.

An animated demonstration and an interactive version to experiment with can be found in [3].

Definition 4 *A continuous map $\theta : [-\alpha; \alpha] \to \mathbb{R}^+$ is called α-probe for $0 < \alpha \leq \pi$ if and only if $\beta > \gamma \geq 0 \Rightarrow \theta(\beta) \leq \theta(\gamma)$ and $\beta < \gamma \leq 0 \Rightarrow \theta(\beta) \leq \theta(\gamma)$.*

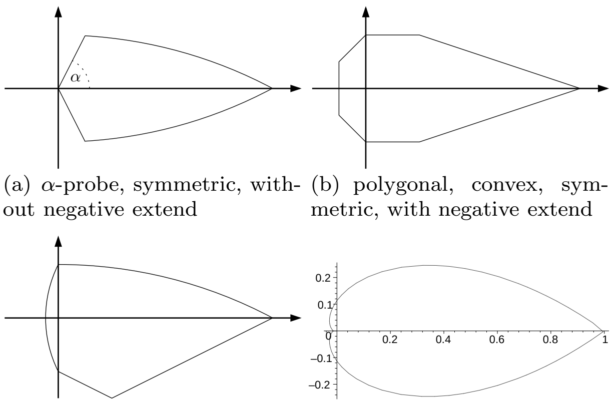

(a) α-probe, symmetric, without negative extend

(b) polygonal, convex, symmetric, with negative extend

(c) convex, asymmetric, with negative extend

(d) $\theta(\alpha) = 4^{-\|\alpha\|}$

Figure 3: Four probes illustrating the variety of possible shapes. All these probes can be described with constant complexity. Probe (a) is assembled from two circular arcs, (b) is a polygon with seven vertices, (c) is a combination of circular arcs and straight segments and (d) is an exponential functions of the angle.

For $\alpha = \pi$, $\theta(\pi) = \theta(-\pi)$ must be true additionally.

An α-probe θ is symmetric *if and only if $\forall\beta \in (0;\alpha] : \theta(\beta) = \theta(-\beta)$.*

An α-probe θ has negative extend *if and only if $\alpha > \frac{\pi}{2}$.*

The name "probe" refers to the shape one obtains from drawing the function θ in polar coordinates with θ as distance from the origin.

Inflating a probe is equivalent to minimizing the following distance function.

Definition 5 *Let θ be an α-probe by definition 4 and $\beta = \angle\overline{pq}, \overline{qr}$. The* probe distance function *for three points p, q, r is defined as*

$$F_{pq}(r) = \begin{cases} \infty & \text{if } \|\alpha\| < \|\beta\| \\ \frac{\|q-r\|}{\theta(\beta)} & \text{otherwise.} \end{cases}$$

Please note that F_{pq} is strictly positive and directed and in general different from F_{qp}.

This definition is valid for any dimension since only the plane spanned by p, q, r is considered and $F_{pq}(r)$ is computed in that plane.

The proof from section 4.1 still applies for arbitrary α-probes with $\alpha \leq \frac{\pi}{2}$ by changing the distance $\|q-r\|$ to $\min\limits_{-\alpha\leq\beta\leq\alpha}(\theta(\beta))$ and $\|q-s\|$ to $\max\limits_{-\alpha\leq\beta\leq\alpha}(\theta(\beta))$.

5 Practical Extensions

5.1 Corners and Endpoints

Sharp corners might become reconstructed with probes with negative extend, first going straight into the corner and then backwards out to a close point (if the apex of the corner is in the sample). Obviously corners with an apex angle of β can only be reconstructed with α-probes with $\alpha \geq \pi - \beta$. Endpoints are the extreme case of corners where the point one came from is the point with minimum distance, i.e. $F_{pq}(p)$ is minimal.

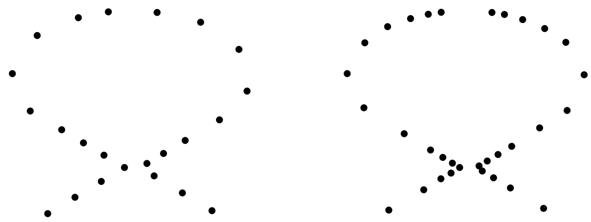

Figure 4: Higher density perceptually indicates disconnectedness. The sample on the left shows nearly equally distributed sample points while the sample on the right uses higher density close to endpoints and corners to emphasize the upcoming gap.

One idea to realize this comes from human perception and should be clear from figure 4. Changing the local density of the sample makes intended gaps much more noticeable. By placing a single additional sample point p very close to an endpoint q, a π-probe aligned at the edge (p, q) with a tiny extend in the backward direction will detect p as the π-neighbor of q. This edge is already part of the recontruction, so the algorithm stops this branch with a vertex of degree one.

Unfortunately an ε-sample is now no longer possible because the medial axis goes through corners, see figure 1. It is also no longer sufficient to guarantee a correct reconstruction from an ε-sample because the distance between two adjacent samples must not be too small if none of them is an endpoint. The latter problem can be solved by switching to (ε, δ)-samples which have the additional condition that for any two samples $p, q \in S$, $\|p - q\| > \delta \cdot \mathrm{lfs}(p)$. This approach is not investigated here, because the necessity for infinite sampling density at corners and intersections remains even for (ε, δ)-samples.

5.2 Intersections

An observation for smooth curves or smooth parts of curves in the vicinity of intersection points is that the angle between segments of properly placed sample points is small while this angle often abruptly increases connecting to a wrong point. Therefore a key idea is to favor small angles over small distances and hope to "tunnel through" a narrow set of wrong points and reach the correct one. This can be realized easily using probes.

A sample as described throughout section 5 always exists. Some construction hints follow but a nice and general rule seems to be hard to find.

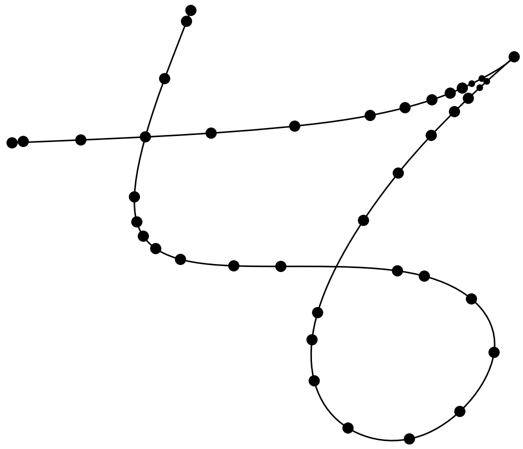

Figure 5: A curve with endpoints, a corner and two intersections. The sample points show the rule of thumb how to sample these features.

For the smooth parts of the curve, an ε-δ-sample suffices. Corners with apex angle β can be sampled loosely like smooth bends without the apex in the sample. If the apex is part of the sample, the distance to the closest point on both edges is given by the negative extend of the used probe. If an intersection point is in the sample, the reconstruction of its vicinity should be simple. Otherwise the adjacent samples should all have roughly the same distance to the intersection point. In the end, one could find the closest pair of sample points and add an even closer point to each endpoint to guarantee correct reconstruction of open curves. Figure 5 gives an intuition of the described rules of thumb.

6 Experimental Results

The Lissajou figure in figure 6(a), given by the parametric form $L(t) \mapsto (\sin 4\pi t, \cos 6\pi t)$, was sampled such that n values were taken equally distributed from $[0; 1)$ and the corresponding points were put into the sample. This was done for several values of n.

The results depend of course on the random input but they show a general behavior which can be reproduced for other inputs of the same size. One possible drawback of the algorithm is that a single failure—a single wrong edge—can lead to a chain reaction for the following edges because they all base on a wrong edge. The experiments show that this disastrous effect is not likely to occur if the sample has at least a certain density.

Obviously one can always create a point pattern around an intersection which will result in a wrong reconstruction independent of the density. Nevertheless there are no additional wrong edges besides close to intersections and no gaps, so the reconstructed topology will be correct if the sample is dense enough. This

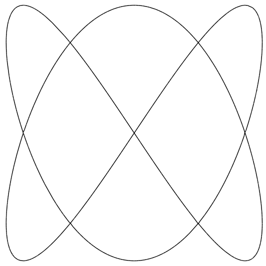

(a) This figure was sampled and reconstructed.

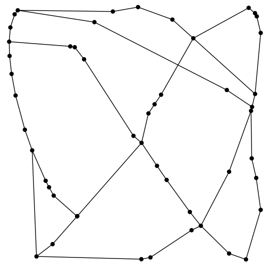

(b) n=50: The sample is overall too sparse.

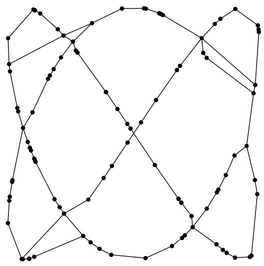

(c) n=100: Some wrong edges appear resulting from sharp bends.

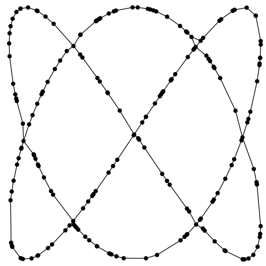

(d) n=200: Almost perfect result except for minor dents and one wrong edge.

Figure 6: Experimental results from random samples.

suggests that the algorithm can also be used successfully in a heuristic approach.

7 Open Problem

The major open problem is the lack of a nice sampling condition. The ε-sampling condition is well-established far beyond the problem of curve reconstruction but it is just not applicable for corners and intersections. Excluding regions around these artifacts and handling them with special cases is possible but definitively not a very nice solution.

Is there a sampling condition, as simple and clear as the ε-sampling condition, which also holds for non-smooth or even self-intersecting curves?

References

[1] T. K. Dey and P. Kumar. A simple provable algorithm for curve reconstruction. In *Proc. 10th ACM-SIAM Sympos. Discrete Algorithms*, pages 893–894, Jan. 1999.

[2] S. Funke and E. A. Ramos. Reconstructing a collection of curves with corners and endpoints. In *Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms*, pages 344–353. Society for Industrial and Applied Mathematics, 2001.

[3] T. Lenz. Reconstructing collections of arbitrary curves. In *SCG '05: Proceedings of the 21st annual symposium on Computational geometry*, pages 366–367. `http://compgeom.poly.edu/acmvideos/socg05video/index.html`, 2005.