




SIGDT: 2D Curve Reconstruction

D. Marin¹  and S. Ohrhallinger¹  and M. Wimmer¹ 

¹Institute of Visual Computing & Human-Centered Technology, TU Wien, Austria




Figure 1: Starting from unstructured points (left), our proximity graph SIGDT (centre) already contains the reconstructed boundary (right).

Abstract

Determining connectivity between points and reconstructing their shape boundaries are long-standing problems in computer graphics. One possible approach to solve these problems is to use a proximity graph. We propose a new proximity graph computed by intersecting the to-date rarely used proximity-based graph called spheres-of-influence graph (SIG) with the Delaunay triangulation (DT). We prove that the resulting graph, which we name SIGDT, contains the piece-wise linear reconstruction for a set of unstructured points in the plane for a sampling condition superseding current bounds and capturing well practical point sets' properties. As an application, we apply a dual of boundary adjustment steps from the CONNECT2D algorithm to remove the redundant edges. We show that the resulting algorithm SIG-CONNECT2D yields the best reconstruction accuracy compared to state-of-the-art algorithms from a recent comprehensive benchmark, and the method offers the potential for further improvements, e.g., for surface reconstruction.

CCS Concepts

• **Computing methodologies** → **Point-based models**;

1. Introduction

Reconstructing a curve based on given samples with no additional information other than their position is a difficult task, considering that no connectivity information is present. As a fundamental problem, with extension to surface reconstruction, it has received a lot

of attention in the field during the last decades. The reconstruction usually implies generating a graph on the input points and filtering/adding edges to recover the connectivity. The resulting shape should interpolate all of the input points, and approximate best the boundary of the shape that the points were sampled from. Ideally,

the reconstruction should be agnostic of the distance between samples and preferably not depend on parameters. However, in practice, this proves to be difficult, especially when multiple types of shapes are considered, such as open curves or multiply connected curves, since a larger distance between samples could mean either a hole in the boundary or unevenly spaced sampling. Hence, we restrict our main method to manifolds and curves with sharp corners.

We introduce a new proximity graph, based on the intersection of the *Spheres-of-influence graph* (*SIG*) and the Delaunay triangulation (*DT*), which we name *SIGDT* and present below. We show its good connectivity property and prove that it contains the piece-wise linear reconstruction of the samples for an enhanced bound of a sampling condition. In order to filter the reconstruction edges from the *SIGDT*, we apply the *inflating* and *sculpting* operations from the *CONNECT2D* algorithm [OM13], yielding a manifold boundary for the input point set.

We present the following three contributions:

- We introduce the graph *SIGDT* by intersecting the *SIG* with the *DT*. *SIGDT* represents connectivity well and is parameter-free.
- We show its good connectivity by proving that it contains the reconstruction edges for an enhanced sampling condition bound that conforms very nicely to point sets in practice.
- As an application, we show manifold curve reconstruction by filtering edges from *SIGDT*, surpassing the state-of-the-art.

The full source-code is publicly available at <https://gitlab.cg.tuwien.ac.at/dmarin/sigconnect2d>.

2. Related work

This section will provide an overview of existing curve reconstruction methods and how they relate to our contribution.

An important category of curve reconstruction methods are explicit methods, where the reconstruction represents an interpolation of the input points. The Delaunay triangulation represents a building block for multiple methods in this category due to its theoretical guarantees that are subject to sampling criteria. One of the first methods to use the Delaunay triangulation was *CRUST* [ABE98], where the ϵ -sampling (based on the local feature size, which incorporates the distance to the medial axis and the distance between samples and is defined below in Subsection 3.1) was also introduced. By choosing a subset of Delaunay edges, the reconstruction is guaranteed to be correct for $\epsilon < 0.252$. Our method similarly starts from the Delaunay triangulation but uses a different set of filters to obtain the final reconstruction of the curve.

Another method to reconstruct a curve was introduced by Dey and Kumar [DK99] – *NN-CRUST*, and it is a proximity-based method, similarly to ours. It uses the edges between nearest neighbours as a starting crust, and for every leaf vertex, adds the shortest edge situated in the half-plane defined by the normal on the edge placed at the leaf vertex. Their results show that proximity-based methods capture the boundary shape, but the sampling limitations of this method (reconstruction is guaranteed for $\epsilon < 1/3$) suggest the possibility of further improvements.

The *NN-CRUST* has been further improved in *HNN-CRUST* [OMW16] to guarantee reconstruction up to $\epsilon < 0.47$. This

method uses a similar idea to the half-plane but places the half-plane's normal as the bisector of the chosen edge.

In *CONNECT2D* [OM13], they use an initial graph BC_0 , representing an approximation of the boundary, which they augment by inflating and sculpting to reconstruct the boundary. The initial boundary graph is a subset of the Delaunay triangulation computed such that minimum boundary length is approximated, all of the vertices are interpolated with a degree of at least 2, and are part of a single connected component. This graph is then processed by inflating, in order to aim to make the curve manifold and contain all vertices on the boundary or inside of it. Candidate triangles are considered for non-manifold vertices and sorted by the increase in total boundary length. Candidate triangles are added to the graph in order to make vertices manifold (i.e. degree 0 or 2) until all of the vertices have become manifold. The resulting graph is processed by sculpting next. Sculpting tries to make the boundary interpolate any vertex that is currently isolated and inside of the boundary. This is done by sorting candidate triangles that are incident to isolated vertices and adding their edges to the graph in order to include the vertices to the boundary. They guarantee reconstruction for $\epsilon < 0.5$ but require a non-uniformity ratio between distances of consecutive samples of $u < 1.609$. We use the *inflating* and *sculpting* operations from their algorithm to filter redundant edges from the *SIGDT*.

GATHANG [DW02] introduced a method specialised in reconstructing data sets with sharp corners. They use the angle and the ratio between edge lengths to filter Delaunay edges for the boundary and their performance is best for this sub-category of curves.

Another category of curve reconstruction methods are implicit methods, where a function is computed over the entire domain of the input and the curve is usually approximated as the zero-set of that function. Important mentions in this type of reconstruction are Signed Distance Functions (SDFs) [HDD*92], Poisson-based reconstructions [KBH06] and radial basis functions [CBC*01]. The signed distance functions compute the signed distance between the points in the plane and the curve and reconstruct the curve as the zero-set of the function. Poisson methods formulate the curve reconstruction as computing the curve whose gradient field of an indicator function is the most similar to the normals of the curve. Radial basis functions try to fit a radial basis function to a signed distance field computed over the input and compute the isoline of this smooth function. Our method rather fits in the explicit category, as the input points are interpolated and no function is computed over the domain.

The spheres-of-influence graph - *SIG* [Tou88] has been introduced as a clustering method since it encodes proximity without the need for any parameters. The *SIG* has been used in implicit surface reconstruction methods [KZ04] to approximate local geodesic distances, as well as a density function over the input points to adapt the algorithm. This is an implicit method that uses the local kernels based on *SIG* to reconstruct the surface. The same approach to use the *SIG* for locally defining an implicit function has been applied to the collision between point clouds [KZ05]. However, we are using *SIG* as the main indicator of proximity over the entire input set and we combine it with additional steps to obtain the reconstruction.

An outlier for the explicit/implicit taxonomy is represented by transport methods. Optimal transport [GCSAD11] reconstructs a

curve starting from the Delaunay triangulation and removing vertices and their incident edges by minimising a global cost. This method is efficient for the reconstruction of noisy curves, a group of input that explicit methods cannot usually correctly reconstruct. Another recent method [CSLV20] tries to solve the curve reconstruction problem as an optimal homologous chain problem. They construct a lexicographic ordering on the simplices of the input to find a minimal chain that bounds the input.

A comprehensive collection of results on multiple types of data sets and different curve reconstructions can be found in the 2D Points Curve Reconstruction Survey and Benchmark [OPP*21]. Results show that most algorithms tend to perform better than their theoretical ϵ -sampling guarantee, but using a graph-based ϵ -sampling measurement (defined in Subsection 3.1) for $\epsilon < 1$, we show that our method performs best for manifold reconstruction in practice.

3. Definitions and Background

This section provides an overview of the main concepts related to curve reconstruction and sampling conditions, most of these having been introduced in the seminal paper [ABE98]. We then describe the spheres-of-influence graph used in our method.

3.1. Sampling of the Curve

We define a set of n points P sampled from a smooth planar curve C . The *medial axis* is defined as the closure of points in \mathbb{R}^2 that are closest to at least two points on the curve C . The local feature size of a point on the curve is defined as the shortest distance from the point to the medial axis. The local feature size is used in combination with the distance between the samples on the curve to define ϵ -sampling as follows: a curve is ϵ -sampled by the point set P if for each point $c \in C$, the ratio between the distance to the closest sample in P and its local feature size is less than ϵ .

For $\epsilon < 1$, the Delaunay triangulation of P is proven to include the reconstruction that best approximates the original shape [ABE98]. However, for sharp angles, where the medial axis touches the curve, such an ϵ -sampling would require infinitely many closely spaced points. Since ϵ -sampling is also used as a measurement for how densely a curve is sampled, in the case of a ground truth in the form of a planar graph, this type of sampling would not result in a meaningful evaluation. Hence, we use an alternate definition of local feature size, particularly defined for planar graphs [lfs95]: the radius of the smallest ball touching two disjoint features (i.e., vertices or edges) of the graph. In order to use this as a base for ϵ -sampling, we divide the distance to the closest sample by this planar local feature size, similarly to above.

3.2. Sphere-of-Influence Graph

The sphere-of-influence graph (*SIG*) was introduced as a clustering method [Tou88]. In the *SIG*, two vertices are connected if the distance between them is less than or equal to the sum of the distance to their respective nearest neighbours.

Visually, the *SIG* can be interpreted as centring a circle at each




Figure 2: Visual representation of the *SIG* connectivity.

vertex whose radius is equal to the distance to its nearest neighbour, i.e., just touching this closest vertex, as can be observed in Figure 2. We connect with an edge all vertices whose circles intersect. This relation encodes spatial proximity without requiring a fixed number of neighbours that the user has to specify, such as the k -neighbourhood. Unlike many other proximity graphs, the *SIG* is not a subset of the *DT* and does not include a triangulation of the input. This motivates the choice of using edges that are part of both graphs as the initial graph.

4. Method

Here we first present the *SIGDT*, show its superior connectivity as a proximity graph, and prove its property of containing the reconstruction. Then, we propose a curve reconstruction algorithm as an application that filters edges of the *SIGDT* with operations from the *CONNECT2D* [OM13] algorithm (described in Section 2) for which pseudo-code is listed as Algorithm 1 and a step-by-step illustration is presented in Figure 4.

We will next provide the definition of the terms used in the algorithm. The boundary is the connected set of edges that includes all edges of the graph G in its interior, e.g., the non-dashed edges in Figure 4f. A non-manifold vertex has more than two incident edges on that boundary, i.e., there the boundary is pinched together, or the spaces defined by its incident edges are exterior, see Figure 4e. An isolated vertex is not included in the boundary, see Figure 4f.

4.1. Boundary-containing Proximity Graph *SIGDT*

Since the *SIG* is not contained in the *DT* but the latter has nice properties such as representing a decomposition of the plane into triangles (which is useful for applications such as curve reconstruction), we design a new proximity graph as the intersection of *SIG* and *DT*, which we name *SIGDT*. This graph combines the advantage of the local proximity offered by the *SIG* with the maximisation of minimum angle triangles provided by the *DT*. Figure 3 shows a visual comparison with the BC_0 proximity graph, which is a superset of the *EMST* constraining vertex degree to ≥ 2 instead of ≥ 1 and is used for *CONNECT2D* curve reconstruction. While *SIGDT* contains all the edges of the reconstruction, BC_0 misses many of them. We show further quantitative comparisons of *SIGDT* with other proximity graphs in the results in Section 5.

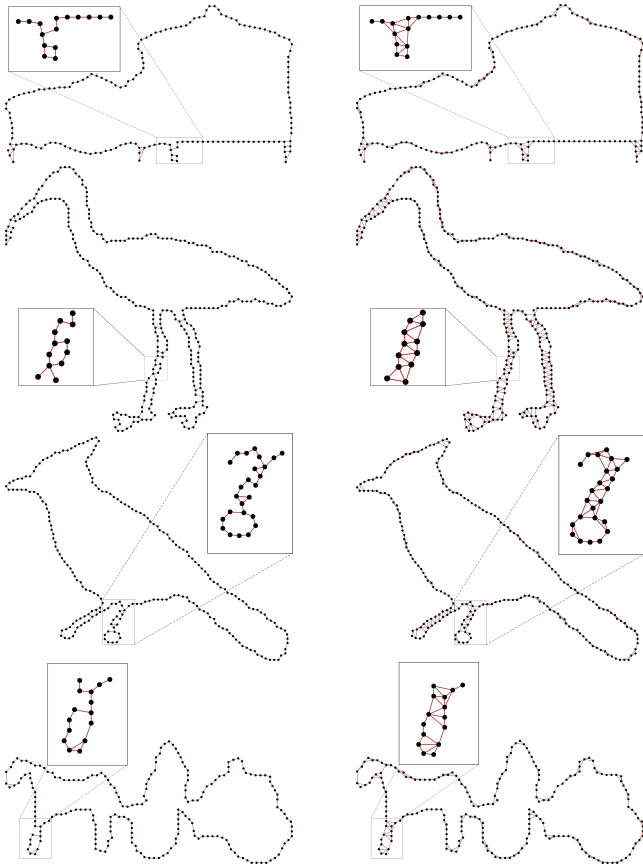


Figure 3: Comparison of BC_0 (left) and $SIGDT$ (right) in terms of encoding proximity. $SIGDT$ manages to include all of the required edges of the reconstruction while BC_0 misses many edges.

Algorithm 1 SIG-Connect2D

Require: Input point set P , Output edge set R

```

1:  $G = \{\}$ 
2: Compute Delaunay triangulation  $DT(P)$ 
3: for  $p \in P$  do
4:   Compute  $NN$  as shortest edge incident to  $p$ 
5: end for
6: for  $p \in P$  do
7:   for  $q \in 1 - ring - neighbourhood(p)$  do
8:     Add edge  $pq$  to  $G$  if  $pq \leq NN(p) + NN(q)$ 
9:   end for
10: end for
11: while  $\exists$  non-manifold vertices in boundary( $G$ ) do
12:    $T = \text{triangles} \in DT$  exterior to and incident to boundary( $G$ )
13:   Add  $t \in T$  that least increases the length of boundary ( $G$ )
14: end while
15:  $G = \text{boundary}(G)$ 
16: while  $\exists$  isolated vertices  $\in DT$  interior to boundary( $G$ ) do
17:    $T = \text{triangles} \in DT$  interior and incident to boundary( $G$ )
18:   Add  $t \in T$  that least increases the length of boundary( $G$ )
19: end while
20:  $R = \text{boundary}(G)$ 

```

We compute $SIGDT$ as follows: In order to determine whether a Delaunay edge is in the SIG , we need to check whether its length is smaller than the sum of the nearest neighbour distance of both its vertices. For that, we first determine the shortest incident edge per vertex in the DT and store its length as its *nearest neighbour distance*. Then, we process all Delaunay edges in the DT and add conforming edges to the SIG definition in order to create the $SIGDT$. This results in $\mathcal{O}(n \log n)$ time complexity.

The $SIGDT$ graph guarantees containing the reconstruction under some sampling condition. In order to prove this, we first have to repeat some definitions.

The reach is defined as the minimum local feature size along an interval between two consecutive samples:

Definition 1 The reach of interval I is $\inf_{p \in I} \text{lfs}(p)$ [Fed59].

Definition 2 A smooth curve C is ρ -sampled by point set P if every point $p \in C$ is closer to a sample than a ρ -fraction of the reach of the interval $I(s_0, s_1)$ of consecutive samples containing it. That is, $\forall p \in I = [s_0, s_1]$ with $s_0, s_1 \in P$: $\|p, s_{[0,1]}\| < \rho \text{reach}(I)$ [OMW16].

Definition 3 The local non-uniformity ratio u is the ratio between the longer and the shorter distance of a sample to its neighbours on the curve: $u = \frac{d_l}{d_s}$ [OM13].

Proof for $\rho < 1, u < 2$: We will now show that $SIGDT$ guarantees to contain the reconstruction of C if it is sampled with $\rho < 1$ and a local non-uniformity ratio of $u < 2$ between distances to samples adjacent on C . A $\rho < 1$ -sampling is equivalent to an $\epsilon < 0.5$ -sampling [OMW16]. This improves on CONNECT2D [OM13], which proves reconstruction for $\epsilon < 0.5$ as well but requires $u < 1.609$, and handles a more relaxed ϵ -sampling than the $\epsilon < 0.47$ for HNN-CRUST [OMW16], although the latter does not require any uniformity. We construct the proof by showing that each edge between consecutive samples is contained in the SIG as well as in the DT .

We need to repeat two statements [ABE98] for our proof:

Corollary 1 A disk centred at a point $p \in C$ with radius at most $\text{lfs}(p)$ intersects C in a topological disk (Corollary 4).

Lemma 1 Any Euclidean disk containing at least two points of a smooth curve in the plane either intersects the curve in a topological disk or contains a point of the medial axis (or both) (Lemma 1).

Now we can prove the following theorem for the $SIGDT$ graph:

Theorem 1 $SIGDT$ contains the reconstruction R of a smooth planar curve C from a set of points P that is sampled with $\rho < 1, u < 2$.

Proof See Figure 5 for illustration. Let s_0, s_1, s_2, s_3 be consecutive samples on C and we want to show that the edge $e(s_1, s_2) \in R$. For this, we need to prove both $e \in SIG$ and $e \in DT$.

1. Prove that $e \in SIG$: The non-uniformity ratio $u < 2$ requires that $\|s_0, s_1\|, \|s_2, s_3\| > \frac{\|s_1, s_2\|}{2}$, yielding $\|s_0, s_1\| + \|s_2, s_3\| > \|s_1, s_2\|$. Thus, the disks centred at s_1, s_2 , with radii $\|s_0, s_1\|, \|s_2, s_3\|$ respectively, overlap. Hence, e is part of SIG , provided that s_0, s_3 are the nearest neighbour samples on C to s_1, s_2 , respectively. This is the case since the disk centred at s_1 with radius $\|s_0, s_1\|$ intersects C in the interval $J[s_0, t]$, with t in

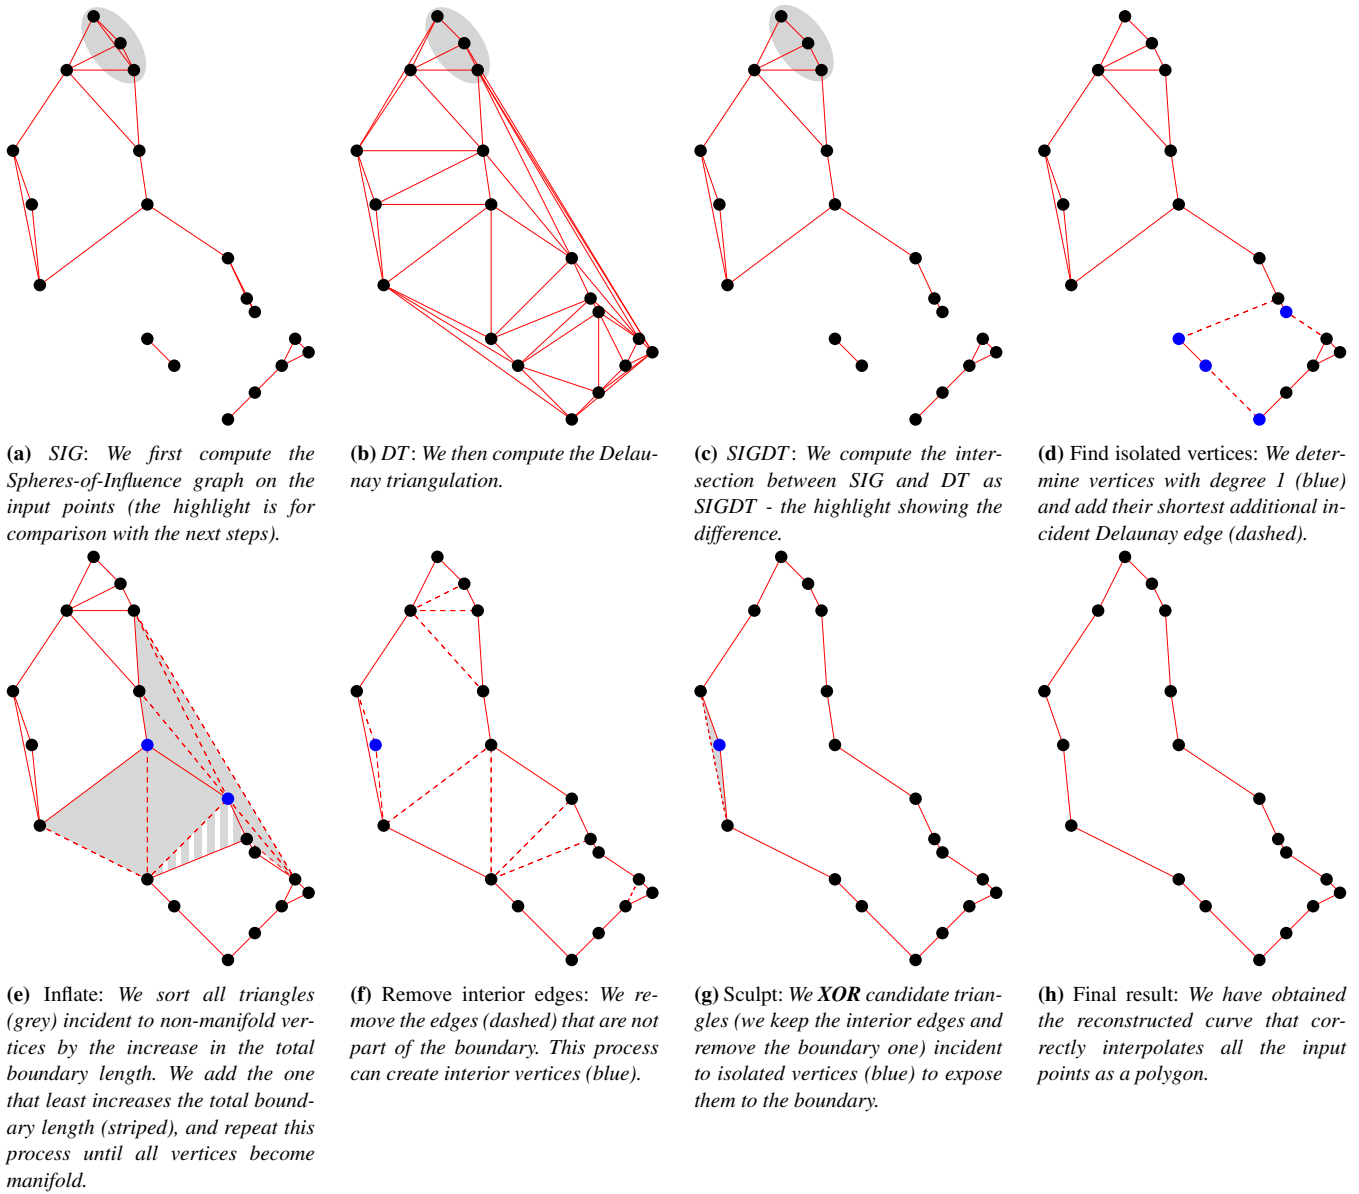


Figure 4: Overview of SIG-CONNECT2D algorithm.

the interval $I[s_1, s_2] \in C$. Thus that disk does not contain any other samples according to Corollary 1. This is proven similarly for s_2, s_3 .

2. Prove that $e \in DT$: For the point $x \in I$ farthest from s_1, s_2 , $\|x, s_{[1|2]}\| \geq \frac{\|e\|}{2}$. Since the curve is ρ -sampled with $\rho < 1$, $\text{reach}(I) > \frac{\|e\|}{2}$ and thus $\text{lfs}(p) > \frac{\|e\|}{2}$ for any $p \in I$. If the smallest disk D including e is empty of other samples, e is a Gabriel edge. Since the Gabriel graph $\subseteq DT$, this would mean that $e \in DT$. We prove D to contain only I by contradiction: Assume a point $q \in C \setminus I$ to exist in D . Then, $C \cap D$ is not a topological disk, and therefore D contains a medial point of C (Lemma 1). Since D has

radius $\frac{\|s_1, s_2\|}{2}$ and contains a medial point, there exists a point $q \in C \cap D$ with $\text{lfs} < \frac{\|e\|}{2}$ which contradicts above $\text{lfs} > \frac{\|e\|}{2}$.

□

Having shown that an edge e between consecutive samples along C under $\rho < 1$ and $u < 2$ is part of both SIG and DT, we have proven that the SIGDT contains the reconstruction under the given sampling conditions.

Corollary 2 Since Theorem 1 [OMW16] proves that any $\varepsilon < r$ -sampling is also a $\rho < r/(1-r)$ -sampling, an $\varepsilon < 0.5$ -sampling is also a $\rho < 1$ -sampling and contains the reconstruction of C if $u < 2$.

The above theorem only guarantees the SIGDT to contain the




Figure 5: The red curve C passes through consecutive samples s_0, s_1, s_2, s_3 and contains the intervals $J[s_0, t]$ for $t \in I[s_1, s_2]$ and $I[s_1, s_2]$. $x \in I$ is the point farthest from the end points of I . The edge e is in SIG because the white disks centred at its end points s_1, s_2 with radii as nearest neighbor distances overlap. The shaded disk D centred at edge e must be empty of samples other than s_1, s_2 to be in DT (by being in the Gabriel graph).

reconstruction edges but may contain additional edges in both its interior and exterior.

4.2. Using SIGDT with CONNECT2D's Dual Operations


We apply two steps from the CONNECT2D algorithm [OM13] to the SIGDT: *inflating* and *sculpting*, in order to improve our reconstruction. These additional steps are summarised below, and they greedily minimise the total edge length of the boundary.

4.2.1. Inflating SIGDT to a Manifold

The boundary subset of SIGDT, named B , contains all vertices either on B or its interior. B may thus not be a manifold as it can be pinched at vertices that have > 2 incident boundary edges. Inflating transforms such non-conforming vertices into manifold ones by selecting incident triangles *exterior* to the boundary and adding their edges to B so that the triangle becomes interior to the boundary. Candidate triangles for all non-conforming vertices are sorted in a priority queue in ascending order by the increase in total boundary length, which is computed by adding the length of new edges and subtracting the length of edges to be removed. We add candidate triangles to non-conforming vertices until they become manifold. However, by adding the edges of new triangles to the graph, some of the edges can become interior to the boundary. Hence, we remove any edge that is not incident to a triangle marked as *outside* (i.e. we remove all edges that are not on the boundary). This procedure is performed in $\mathcal{O}(n \log n)$ time complexity and guarantees a manifold boundary B' as its result. More details on the exact implementation and proofs of the theoretical guarantee can be found here [OM13]. The inflating procedure is illustrated in Figure 6.

4.2.2. Sculpting the Manifold to Interpolate Interior Vertices

The manifold boundary B' resulting from inflating contains all vertices either on B' or interior to it. These isolated interior vertices




(a) We identify the incident exterior triangle to a degree ≥ 2 vertex with the least boundary length change.

(b) Adding its edges to the boundary creates a new degree ≥ 2 vertex, so we add another such triangle.

(c) Removing the interior edges results in a manifold boundary.

Figure 6: Step-by-step inflating procedure on a close-up to make the curve manifold.



(a) We first locate isolated points such as in this example.

(b) We identify its incident interior triangle with the least boundary length change.

(c) We XOR the triangle's edges to the boundary, interpolating that interior point.

Figure 7: Step-by-step sculpting procedure on a close-up of a manifold curve to interpolate interior points.

have to be connected to the boundary so that the reconstruction can interpolate all the points. Triangles incident to vertices *interior* to B' are sorted by the same boundary length increase criterion as for inflating. When a candidate triangle is added to B' , we XOR that triangle's edges with B' (i.e. remove the triangle edge that is already part of the graph and add the other edges that are not yet part of the graph). This step does not increase the overall complexity of the algorithm, being performed in $\mathcal{O}(n \log n)$ time. This step exposes the interior vertices to the boundary. This fails if the DT does not contain a Hamiltonian cycle. The details on the procedure and guarantees can be found here [OM13]. The sculpting process is illustrated in Figure 7.

4.2.3. Eliminate Leaf Vertices (optional)

In the cases where C is not sampled as densely as required, artefacts such as leaf vertices may appear. In our experiments, we found that this does not affect results in general but eliminating these further improves reconstruction quality for point sets with sharp corners. We apply this optional step to the SIGDT before *inflating*:

We increase the degree of the leaf nodes to two by adding their shortest incident Delaunay edge to SIGDT, forming SIGDT2, which has vertex degree ≥ 2 everywhere. This step loops over all the vertices and finds the shortest incident edge to each leaf vertex that is not part of the graph yet. This takes $\mathcal{O}(kn)$ operations, as-

suming a constant degree k of vertices, for the average DT , ignoring contrived cases which do not arise often in practice.

As an overview, we have described SIG-CONNECT2D as a SIG -based curve reconstruction method that uses an unstructured point set as input and reconstructs the boundary of the shape that the points have been sampled from. SIG-CONNECT2D starts with the $SIGDT$ graph, and then applies *inflating* and *sculpting* on it as in the CONNECT2D algorithm - an overview of this algorithm is presented in Figure 4. The next section presents our results compared to state-of-the-art curve reconstruction algorithms.

5. Results

We have tested our proposed method, SIG-CONNECT2D, against 15 state-of-the-art curve reconstruction algorithms (see names in Figure 8) using the 2D Points Curve Reconstruction Survey and Benchmark [OPP*21], where these are referenced together with source code and the data sets used for our evaluations below. Note that OPTIMALTRANSPORT has been eliminated from their evaluation since the input it aims to reconstruct is dense, with a high percentage of outliers and noise, and cannot exactly reconstruct clean, sparse inputs (an example of how it fails can be seen in Table 1), being also highly dependent on the number of iterations. We have then analysed results for the reconstruction of manifold curves, curves with sharp corners, and a subset of well-sampled manifold curves in terms of exact reconstruction, and examples can be seen in Figure 20. Furthermore, we analysed our method in terms of the root mean square error (RMSE) to the ground-truth curve for noise-free data sets, noisy data sets, and data sets with outliers to form a thorough evaluation. Also, we compute the overlap of sets of edges of the proximity graphs SIG , DT , $SIGDT$ and reconstruction for a large data set.

Manifold Curves: We have compared our results against the above-mentioned reconstruction algorithms on 1257 noise-free point sets. These data sets represent a subset of the original benchmark data set since we have chosen to only use ground-truth data sets that interpolate all input points. Our algorithm SIG-CONNECT2D shows the best accuracy (91.5% compared to second-best CONNECT2D with 90.3%). Figure 8 shows the improved reconstruction as a visual comparison on manifold data sets, numbers are given in Table 2. An example of a point set fed to all algorithms is shown in Table 1 where only our SIG-CONNECT2D reconstructs it correctly.

Proximity Graphs Overlap: Using the same data sets, we have computed the average percentage of SIG edges that are also in the DT and vice-versa, as well as the average percentage of $SIGDT$ edges that are part of the reconstruction and vice versa. The results show that SIG is usually mostly contained in the DT - 90.9%, especially considering that the nearest neighbour graph is contained in both. However, as expected, DT contains more edges that are not part of SIG - only 47.2% of DT edges are in SIG . Furthermore, in practice, even without imposing the sampling and non-uniformity criteria, the reconstruction is contained in $SIGDT$ in almost all cases (99.9%), indicating that our sampling condition covers practical point sets extremely well. This represents the best overlap among tested proximity graphs - BC_0 achieves to contain,

on average, 99.6% of the ground-truth edges, while kNN graphs, for $k \in \{2, 3, 4\}$, achieve at most 98.6% ($2NN$ - 98.4%, $2NN$ - 98.5% and $4NN$ - 98.6%). However, $SIGDT$ usually has more edges than the reconstruction - only 76.8% of $SIGDT$ edges are part of the correct reconstruction. These results show that the $SIGDT$ is a good indicator of the proximity of the graph, as it misses almost no edges.




Figure 8: Reconstruction of manifold curves.

Sharp Corner Curves: We have tested our method on a data set consisting of 47 input sets, comparing it against the same other curve reconstruction methods as above. The best results are achieved by GATHANG [DW02], at 80.9% accuracy of the exact reconstruction, followed by our method at 70.2% - the complete results are presented in Figure 9 and Table 2. However, these results are expected since GATHANG is specialised for sharp-corner reconstruction, but performs worse in the general case of manifold curves as seen above.




Figure 9: Reconstruction of curves with sharp corners.

Open and Multiple Curves: Our method is guaranteed to output a manifold reconstruction of the output through the usage of inflating and sculpting [OM13]. For this reason, SIG-CONNECT2D is not suitable for such types of input. However, we provide the symmetric difference in area (computed using BOOST's `boost_sym_difference`) between the output of different algorithms and the correct output. Our method interpolates the input

SIG-CONNECT2D		CONNECT2D		LENZ		OPTIMALTRANSPORT	
CRUST		CCRUST		HNNCRUST		NNCRUST	
DISCUR		VICUR		CRAWL		PEEL	
GATHANG		GATHANI		FITCONNECT		STRETCHDENOISE	

Table 1: The resulting reconstructions for our algorithm compared to the other 15 state-of-the-art algorithms [OPP*21] on manifold curve input. Our SIG-CONNECT2D is the only one to correctly reconstruct this point set.

points and does not achieve exact open curves or multiple curves as expected, but the results are still similar to the expected output, as visible in the symmetric difference results in Figure 10 and Figure 11, and examples are presented in Figure 21.

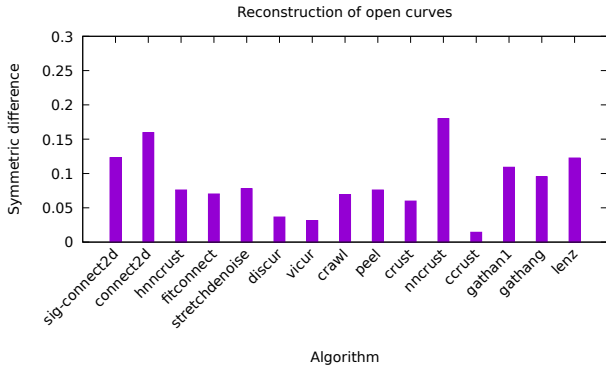


Figure 10: Symmetric area difference for open curves.




Figure 11: Symmetric difference of area for multiple curves.




Figure 12: Reconstruction accuracy for graph-based $\epsilon < 1$.

Well-Sampled Manifold Curves: Since some of the ground-truth curves in the data set are sparsely sampled, and therefore cannot be reconstructed well by any algorithm, we have selected a subset of the 1257 data sets from the manifold curve test set. We

filter 1183 data sets ϵ -sampled with $\epsilon < 1$ based on local feature size computed on graphs as explained in Section 3 and repeat the above comparison on manifold curves. Our algorithm performs best at 95.8%, followed by the original CONNECT2D at 95.0%, showing that the SIG captures the connectivity better for well-sampled curves, and comes quite close to reconstructing all curves sampled with graph-based $\epsilon < 1$ in practice. Results are presented in Figure 12 and in Table 2.




Figure 13: Average runtime of manifold curve reconstruction.

Runtime: We have specified the time complexity of all steps of our method in the respective descriptions, and their upper bound is $\mathcal{O}(n \log n)$ in terms of the n input points - typical for the Delaunay-based curve reconstruction algorithms. The average runtime of each algorithm run on the complete data set of 1257 point sets on an Intel Core i7-7700HQ processor is presented in Figure 13 and in Table 2. The empirical results (≈ 1 ms per point set, with an average of 260 points) confirm the theoretical bounds on the time complexity and are in line with the other methods except fastest NN-CRUST. We have also tested on a large point set with 9991 points, which takes 75ms (Table 2), indicating runtime is almost linear.

RMS error: We have tested how closely our reconstruction approximates a cubic Bézier curve by sampling it with different ϵ values and computing the RMS error between the reconstructed curve and the original (see Figure 14). Our algorithm performs similarly to the majority of evaluated algorithms. Furthermore, for the same setup of ϵ -sampling a cubic Bézier curve with $\epsilon = 0.1, 0.2, 0.3, 0.4, 0.5$, we have tested our algorithm on 20 differently generated point sets for each ϵ value by varying the starting sample on the curve. This shows that the theoretical guarantee of the SIGDT includes the accurate reconstruction in all cases, even without constraining u .

Noise: Even if our method is designed for non-noisy input, we have tested its reliability against noise by computing the RMS error against the ground truth. We have added uniform noise to some of the input curves and run the reconstruction algorithms. These results are visible in Figure 16 and indicate that our method performs similarly to CONNECT2D, as expected. Another way of testing the resilience to noise was by adding IFS noise on samples along a cubic Bézier curve. The results are competitive, as shown in Figure 17. We also present the output of running the algorithm on noisy curves

Algorithm	Manifold	Sharp	Open	Multiple	$\epsilon < 1$	Manifold runtime (ms)	Large data set runtime (ms)
SIG-CONNECT2D	91.5%	70.2%	0.0%	3.7%	95.8%	0.9	75
CONNECT2D	90.3%	65.9%	0.0%	0.0%	95.0%	0.5	99
HNNCRUST	64.5%	14.8%	43.4%	53.7%	67.2%	0.6	49
FITCONNECT	67.7%	10.6%	8.6%	22.2%	71.0%	237.1	-
STRETCHDENOISE	64.7%	11.1%	9.5%	24.0%	68.0%	207.4	-
DISCUR	49.0%	17.0%	39.1%	46.2%	50.7%	279.7	-
VICUR	46.2%	21.2%	52.1%	46.2%	47.6%	357.5	-
CRAWL	75.3%	10.6%	21.7%	40.7%	78.2%	0.4	182
PEEL	64.7%	14.8%	43.4%	57.4%	67.3%	2.8	2455
CRUST	71.3%	23.4%	43.4%	38.8%	74.3%	0.6	39
NNCRUST	39.3%	10.6%	8.6%	12.9%	41.3%	0.1	11
CCRUST	53.1%	0.0%	30.4%	22.2%	55.7%	0.8	796
GATHAN1	62.4%	44.6%	13.0%	24.0%	65.3%	0.3	18
GATHANG	63.2%	80.8%	21.7%	35.1%	65.7%	1.9	293
LENZ	2.8%	8.8%	0.0%	0.0%	3.0%	8.2	12672
OPTIMALTRANSPORT	0.0%	0.0%	0.0%	0.0%	0.0%	71.2	542

Table 2: Results for precision of exact reconstruction of curves with different characteristics and average runtime in milliseconds, as well as the runtime for a large point set with 9991 points, for which not all algorithms managed to produce an output.




Figure 14: RMSE of reconstructions for varying ϵ -sampling.

in Figure 15, which achieves, as expected, an interpolation of the input points.

Outliers: We have tested our algorithm's reliability when outliers are present by adding a percentage of outliers to some of our input curves. The results are in line with the majority of algorithms and are displayed in Figure 18.

Limitations: We present in Figure 22 some of the cases where our algorithm fails to produce the exact reconstruction of the input data. However, most of the points of failure are represented by multiple components and non-uniform sampling that go beyond the theoretical limits of our method. These cause the algorithm to try to create a boundary interpolating all components together or to fall into local minima when geodesically far samples become geometrically close. However, for non-manifold curves, we provide an analysis of the symmetric difference of area between the results of various algorithms and the ground truth in Figure 19. Our method




Figure 15: Reconstruction of data sets perturbed with uniform noise as a percentage of the bounding box diagonal. The left data set uses 0.01% uniform noise, and the original shape is correctly reconstructed, while the right one uses 0.03%, thus failing to recreate the ground truth as the sampling becomes too sparse.

fails to produce exact reconstruction of such curves, but the results are close to the original curve, as illustrated in Figure 21.

6. Conclusion and Future Work

We propose a new proximity graph in 2D, $SIGDT = SIG \cap DT$, and show that it better captures connectivity between points than other proximity graphs, and does so without requiring a specific number of neighbours as a parameter, as kNN would. We prove that $SIGDT$ contains the reconstruction for planar curves for some enhanced sampling condition. Together with filtering steps for redundant edges from an existing method, our method SIG-CONNECT2D correctly reconstructs the manifold boundary of the input set in

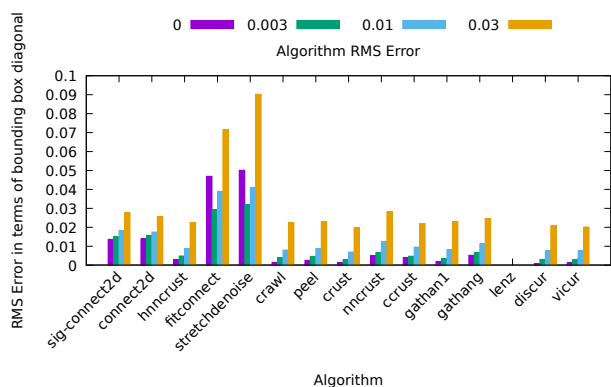


Figure 16: RMSE of reconstructed curves from inputs contaminated with uniform noise.




Figure 17: RMSE of curves from inputs with lfs-based noise.

more cases than the state-of-the-art, and reconstructs almost all well-sampled point sets. As current results are promising, they encourage further improvements in this approach. Hence, our future work includes:

- Extending SIG-CONNECT2D to robustly reconstruct multiple curves, open curves, and non-manifold inputs;
- Extending SIGDT to 3D and SIG-CONNECT2D to surface reconstruction.

Acknowledgements

This work has been partially funded by the Austrian Science Fund (FWF) project no. P32418-N31 and by the Wiener Wissenschafts-, Forschungs- und Technologiefonds (WWTF) project ICT19-009.

References

[ABE98] AMENTA N., BERN M., EPPSTEIN D.: The crust and the beta-skeleton: Combinatorial curve reconstruction. *Graphical Models and Image Processing* 60 (01 1998), 125–135. 2, 3, 4

[CBC*01] CARR J. C., BEATSON R. K., CHERRIE J. B., MITCHELL T. J., FRIGHT W. R., MCCALLUM B. C., EVANS T. R.: Reconstruction and representation of 3d objects with radial basis functions. Association for Computing Machinery. 2

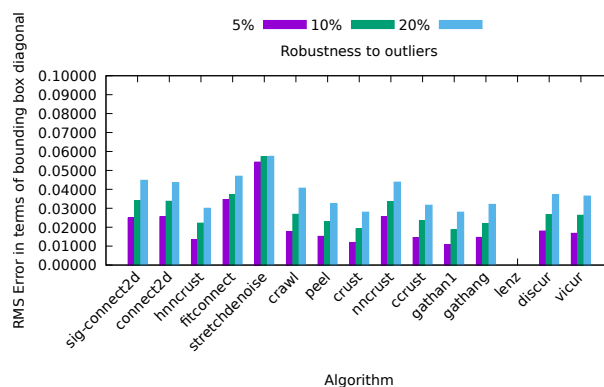


Figure 18: Reconstruction of curves with varying outlier share.

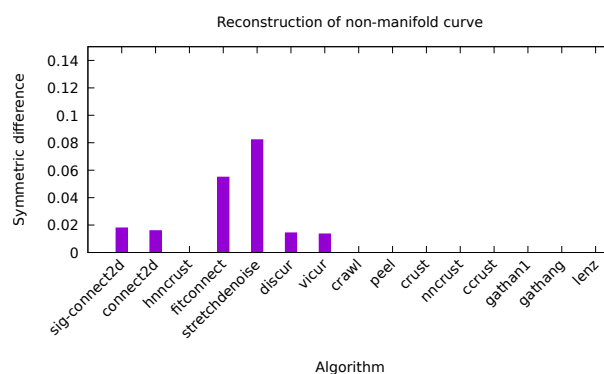


Figure 19: Symmetric area difference for non-manifold curves.

[CSLV20] COHEN-STEINER D., LIEUTIER A., VUILLAMY J.: Lexicographic optimal homologous chains and applications to point cloud triangulations. In *SoCG* (2020). 3

[DK99] DEY T. K., KUMAR P.: A simple provable algorithm for curve reconstruction. In *Proceedings of the Tenth Annual ACM-SIAM Symposium on Discrete Algorithms* (USA, 1999), SODA '99, Society for Industrial and Applied Mathematics, p. 893–894. 2

[DW02] DEY T., WENGER R.: Fast reconstruction of curves with sharp corners. *Int. J. Comput. Geometry Appl.* 12 (10 2002), 353–400. 2, 7

[Fed59] FEDERER H.: Curvature measures. *Transactions of the American Mathematical Society* 93, 3 (1959), 418–491. 4

[GCSAD11] GOES F., COHEN-STEINER D., ALLIEZ P., DESBRUN M.: An optimal transport approach to robust reconstruction and simplification of 2d shapes. *Comp. Graph. Forum* 30 (08 2011), 1593–1602. 2

[HDD*92] HOPPE H., DEROSE T., DUCHAMP T., MCDONALD J., STUETZLE W.: Surface reconstruction from unorganized point clouds. 2

[KBH06] KAZHDAN M., BOLITHO M., HOPPE H.: Poisson Surface Reconstruction. In *Symposium on Geometry Processing* (2006), Sheffer A., Polthier K., (Eds.), The Eurographics Association. 2

[KZ04] KLEIN J., ZACHMANN G.: Point cloud surfaces using geometric proximity graphs. 2

[KZ05] KLEIN J., ZACHMANN G.: Interpolation search for point cloud intersection. pp. 163–170. 2

[lfs95] A delaunay refinement algorithm for quality 2-dimensional mesh generation. *Journal of Algorithms* 18, 3 (1995), 548–585. 3

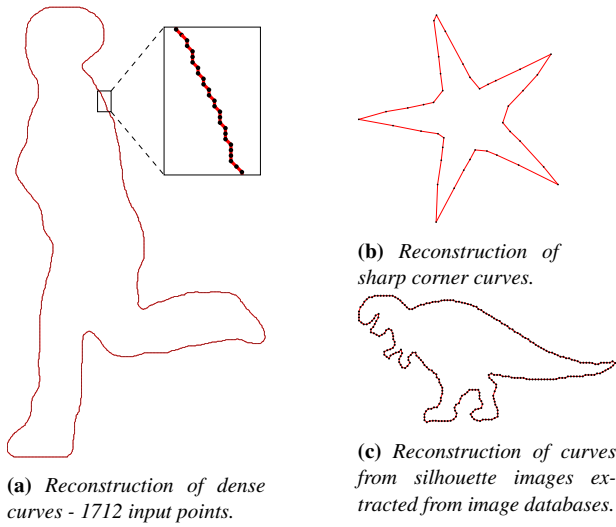


Figure 20: Reconstruction of different types of curves.

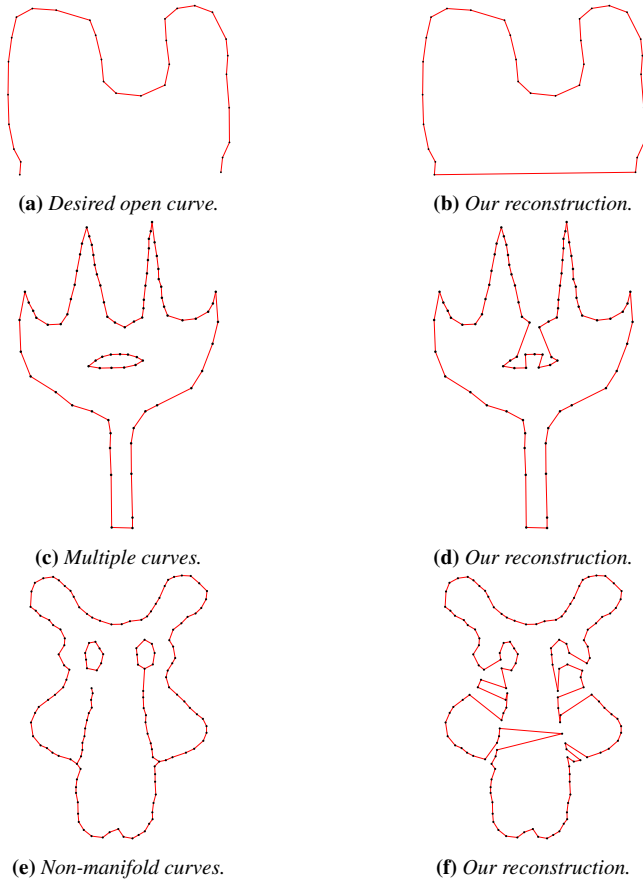


Figure 21: Reconstruction of open, multiple and non-manifold curves. Our method is not designed for these categories since we expect the output to be a manifold curve that interpolates all the input points. However, the results are similar to the expected output.

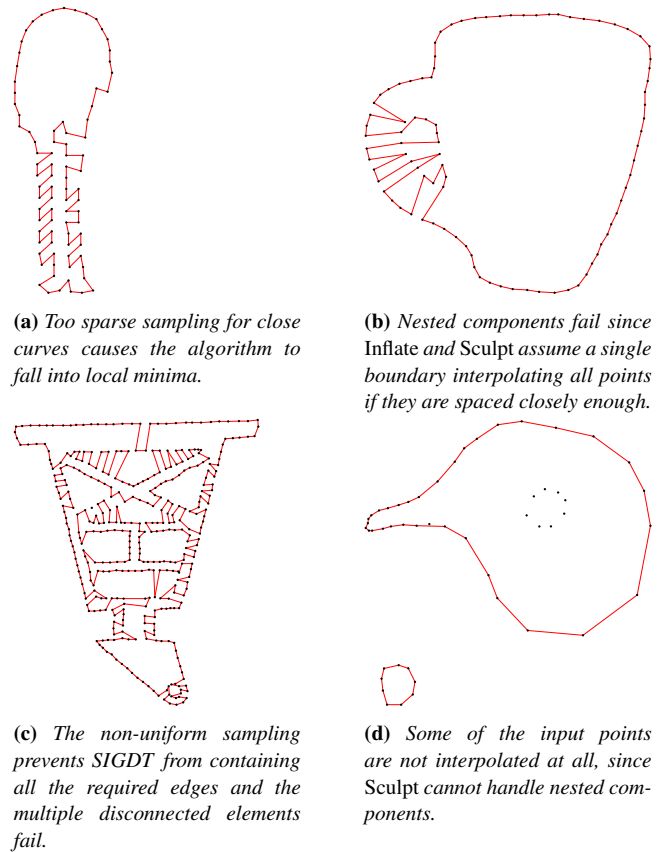


Figure 22: Failure cases for SIG-CONNECT2D.

[OM13] OHRHALLINGER S., MUDUR S.: An Efficient Algorithm for Determining an Aesthetic Shape Connecting Unorganized 2D Points. *Computer Graphics Forum* (2013). 2, 3, 4, 6, 7

[OMW16] OHRHALLINGER S., MITCHELL S., WIMMER M.: Curve reconstruction with many fewer samples. *Computer Graphics Forum* 35 (08 2016), 167–176. 2, 4, 5

[OPP*21] OHRHALLINGER S., PEETHAMBARAN J., PARAKKAT A. D., DEY T. K., MUTHUGANAPATHY R.: 2d points curve reconstruction survey and benchmark. In *Computer Graphics Forum* (2021), vol. 40, Wiley Online Library, pp. 611–632. 3, 7, 8

[Tou88] TOUSSAINT G. T.: A graph-theoretical primal sketch. In *Machine Intelligence and Pattern Recognition*, vol. 6. Elsevier, 1988, pp. 229–260. 2, 3