



Region reconstruction with the sphere-of-influence diagram*

Luiz Henrique de Figueiredo ^{a,*}, Afonso Paiva ^b

^aIMPA, Rio de Janeiro, Brazil

^bICMC, Universidade de São Paulo, São Carlos, Brazil

ARTICLE INFO

Article history:

Received May 9, 2022

Revised July 10, 2022

Final version August 1, 2022

Keywords: shape reconstruction, point samples, computational morphology

ABSTRACT

We describe a simple method for reconstructing regions in the plane from well-distributed point samples. For that, we introduce the sphere-of-influence diagram, a planar diagram extracted from the Delaunay triangulation and the sphere-of-influence graph. The sphere-of-influence diagram is simple to understand and to implement, and supports an intuitive scaling parameter to handle variations in the distribution of samples. We report experiments reconstructing regions from point clouds of varying shape and density. We compare our results with those obtained by α -shapes and by CT-shapes.

© 2022 Elsevier B.V. All rights reserved.

1. Introduction

The reconstruction of shapes from point samples is an important task in pattern recognition, computer vision, and geometry processing. We consider here the reconstruction of a region in the plane from a well-distributed point sample. When such a sample is presented to a human as a dot pattern using small disks, the shape of the region is immediately perceived (see Fig. 1). The algorithmic challenge is that the computer cannot see the dot pattern, only the coordinates of the points.

Region reconstruction from samples is an ill-posed problem. While humans easily perceive a shape from a dot pattern, they do not always agree on the perceived shape. Region reconstruction is a hard problem because there is no clear notion of optimal shape that can be used as a proxy for perception [1]. Moreover, in some formulations it is an NP-hard problem [2]. Most methods for region reconstruction (including ours) start from the Delaunay triangulation of the sample points and sculpt the boundaries of the region by removing edges and faces, an approach known as Delaunay filtering.

Toussaint [3] coined the term *computational morphology* for the use of geometric algorithms to extract shape information and

introduced the sphere-of-influence graph [4, 5] as a tool for computational morphology. He argued that the sphere-of-influence graph is markedly different from other proximity graphs because it captures shapes naturally, identifying connected components and holes without human intervention or tuning of parameters.

In this paper, we introduce the *sphere-of-influence diagram* of a set of points in the plane: it is the planar diagram induced by intersecting the sphere-of-influence graph with the Delaunay triangulation. Surprisingly, to the best of our knowledge this natural diagram has not been studied or used before. We report experiments that show how the sphere-of-influence diagram of the samples reconstructs the boundaries of a region from a well-distributed point sample. We compare our results with those obtained by the widely used α -shapes [6] and also by the recent CT-shapes [7].

The sphere-of-influence diagram introduced here is simple to understand and to implement. It reconstructs regions from spatially uniformly distributed point samples without intervention. It handles samples that are less well distributed by tuning an intuitive scaling parameter $\mu \geq 1$. Tuning μ closes false holes, fixes topological irregularities, and also improves the smoothness of the reconstructed boundaries. When needed, choosing a suitable μ for the sphere-of-influence diagram is easier than choosing a suitable α for α -shapes because $\mu = 1$ gives a good starting point and the range of possible values of μ is much smaller than the range of possible α .

*Dedicated to the memory of Godfried Toussaint.

*Corresponding author: Luiz Henrique de Figueiredo (lhf@impa.br)

2. Related work

The seminal work on shape reconstruction was done in the 1980s. Edelsbrunner, Kirkpatrick, and Seidel [6] proposed the family of α -shapes, which are widely used. Kirkpatrick and Radke proposed β -skeletons [8]. Radke [9] discussed these and other proximity graphs, including a variant of the sphere-of-influence graph that uses a parameter for scaling nearest neighbor disks, which we have adopted here. Toussaint organized a volume on computational morphology [3] and introduced the sphere-of-influence graph [4].

Medeiros et al. [10] restricted ball-pivoting to the plane to compute solid α -shapes directly without computing a Delaunay triangulation. Ball-pivoting traverses the sample points using a ball of fixed radius. They also discussed how to adapt the radius to the density of the samples. Klein and Zachmann [11] proposed extensions of the sphere-of-influence graph to reconstruct implicit surfaces from noisy point clouds. They used the k th nearest neighbor instead of the standard first nearest neighbor. They also proposed anisotropic sphere-of-influence graphs which use ellipsoids instead of spheres. Vital Brazil and Figueiredo [12] used radial basis functions with Gaussian kernels to find an approximation of a fuzzy membership function for a region in the plane from a noisy point cloud, without using proximity graphs.

Motivated by geographic applications, Galton and Duckham [13] discussed the shape of a set of points in the plane attempting to characterize the region they occupy. They introduced the swinging arm algorithm, a generalization of the gift wrapping algorithm for computing convex hulls [14] that uses a line segment of fixed length instead of an infinite half-line. A similar approach was proposed by Jarvis [15].

Later work on region reconstruction focused on Delaunay filtering. Duckham et al. [16] introduced χ -shapes, a single simple polygon that characterizes the shape of a set of points in the plane. Their algorithm repeatedly removes long edges from the Delaunay triangulation of the points until no long edges remain, subject to the regularity constraint that the external boundary of the resulting triangulation is a simple polygon. Their algorithm is controlled by a user-supplied length parameter and does not handle multiple components or holes. Zhong and Duckham [17] extended χ -shapes to χ -outlines, which can handle multiple components and outliers from noisy data. Zhong and Duckham [18] described an incremental algorithm for computing χ -shapes that supports insertions and deletions, and handles streaming individual points and multiple point sets.

Peethambaran and Muthuganapathy [19] combined a relaxed Gabriel graph with a structural characterization of holes to reconstruct regions having a single connected component. Methirumangalath et al. [20] presented a unified reconstruction method that handles both boundary samples and dot patterns. External edges are successively removed subject to empty circle and regularity constraints. The resulting boundary is a simple closed curve. Methirumangalath et al. [21] described how to detect holes using empty circle criteria and applied that to reconstruct external and internal boundaries. The recent work of Thayyil et al. [7, 22] handles region reconstruction from both boundary samples and dot patterns and includes extensive comparative evaluations of existing methods.

3. Problem statement and background

Input. The input data is a set S of n points sampling a bounded region in the plane. The region may have several connected components and the components may have holes. Each component has one external boundary and possibly also internal boundaries that define holes. All boundaries are assumed to be piecewise smooth. The sample points are assumed to be well distributed inside the region. Spatially uniform distributions, such as created by blue noise sampling, are ideal input. The boundaries are not assumed to be sampled, except by accident. The sample points are given by their Cartesian coordinates.

Output. To reconstruct the region, we seek a description the external and internal boundaries of its connected components by simple closed polygonal lines whose vertices come from the point sample. Every sample point is inside exactly one polygon.

Spheres of influence. To each sample point p , we associate the disk centered at p whose radius $r(p)$ is the distance from p to the nearest sample point: $r(p) = \min\{\text{dist}(p, q) : q \in S, q \neq p\}$, where dist is Euclidean distance. We call this disk the *nearest neighbor disk* of p and its radius the *nearest neighbor radius* of p . Toussaint's insight [4] was that the nearest neighbor disk of a point is a measure of its influence in the sample, and two points exert mutual influence when their disks intersect. This insight led him to call nearest neighbor disks as *spheres of influence* and to define three related structures that reflect mutual influence: the sphere-of-influence hull, shape, and graph, which we describe below. Radke [9] introduced a parameter $\mu \geq 1$ that controls influence by scaling nearest neighbor disks. The *scaled nearest neighbor radius* of a sample point p is $R(p) = \mu r(p)$. We use this parameter μ to handle samples that are less well distributed than spatially uniformly distributed samples. In what follows, all nearest neighbor disks and radii are implicitly scaled by μ . Toussaint's remarks when $\mu = 1$ apply equally well when $\mu > 1$.

Sphere-of-influence hull. The union of all nearest neighbor disks is the *sphere-of-influence hull* of the points. When the sample is well distributed, the sphere-of-influence hull with $\mu = 1$ is a fairly good outer approximation of the region [4, 9] (see Fig. 1). This observation by Toussaint is the underlying motivation for this paper. However, computing the boundary of the sphere-of-influence hull is not simple: it requires identifying and intersecting boundary circles, and identifying and ordering boundary arcs. These are complicated and numerically delicate geometric tasks. The boundary has size $O(n)$ and, in principle, can be found in time $O(n \log n)$ using the Laguerre Voronoi diagram of the points [23] (but we know of no actual implementation).

The change on the sphere-of-influence hull as μ varies is reminiscent of dilation in mathematical morphology and filtration in persistent homology for topological data analysis. However, in both these examples all disks have the same radius whereas the change on the sphere-of-influence hull as μ varies is a non-uniform dilation that respects local sample density.

Sphere-of-influence graph. The sphere-of-influence graph describes mutual influence across the whole sample. The vertices are the sample points. Two vertices u and v define an edge




Fig. 1. From left to right: point sample ($n = 258$), sphere-of-influence hull, Delaunay triangulation, sphere-of-influence graph, sphere-of-influence diagram, and region boundaries extracted from sphere-of-influence diagram (all computed with $\mu = 1$).

when their scaled nearest neighbor disks intersect, that is, when $\text{dist}(u, v) \leq R(u) + R(v)$, where dist is Euclidean distance and R is the scaled nearest neighbor radius. Toussaint [4] noticed that, when $\mu = 1$, this simple notion of proximity captures the local structure of the data. In particular, like the hull, the sphere-of-influence graph automatically identifies the connected components of the region. Other proximity graphs are either too sparse (like the nearest neighbor graph) or too connected (like the Euclidean minimum spanning tree and all proximity graphs that contain it, such as the Delaunay triangulation). Moreover, when $\mu = 1$, the sphere-of-influence graph has size $O(n)$ and can be found in optimal time $O(n \log n)$ [5]. When μ is sufficiently large, the sphere-of-influence graph is the complete graph.

Sphere-of-influence diagram. Even when $\mu = 1$, the sphere-of-influence graph is typically not planar, despite having linear size: it tends to contain small cliques and so edges that cross (see Fig. 1). Thus, the sphere-of-influence graph cannot be used directly to extract boundaries because it has no natural faces. On the other hand, the Delaunay triangulation of the points is planar, but it joins different components, closes holes, and insists on a single convex boundary enclosing all points.

We introduce here the *sphere-of-influence diagram* as the planar diagram induced by intersecting the sphere-of-influence graph with the Delaunay triangulation. The vertices are the sample points. The edges of the sphere-of-influence diagram are the Delaunay edges that are also in the sphere-of-influence graph. The faces of the sphere-of-influence diagram are the Delaunay faces whose edges are all in the sphere-of-influence graph. Thus, by construction, the sphere-of-influence diagram is a simplicial complex contained in the Delaunay triangulation and so is planar and has size $O(n)$. When μ is sufficiently large, the sphere-of-influence diagram coincides with the Delaunay triangulation.

As a simplicial complex contained in the Delaunay triangulation, the sphere-of-influence diagram is composed of vertices which are points, edges which are line segments, and faces which are triangles (see Fig. 1). There are three types of edges in the diagram: *internal* edges, which belong to two faces, *external* edges, which belong to only one face, and *free* edges, which belong to no face. The external edges compose the boundaries of the sphere-of-influence diagram. The free edges play no role.

The sphere-of-influence diagram provides an inner approx-

imation of the region and serves much the same purpose as the sphere-of-influence hull, but it is much easier to compute. Our main empirical observation is that, already for $\mu = 1$, the sphere-of-influence diagram of a spatially uniformly distributed sample is a regular simplicial complex, and so has well-defined boundaries that are easily extracted using standard topological procedures. Moreover, we also observe empirically that we can reconstruct regions from samples that are less well distributed by tuning $\mu > 1$. See the results reported in §5.

4. Method

Finding the sphere-of-influence diagram. The sphere-of-influence diagram is extracted easily in linear time from the Delaunay triangulation, which can be found in optimal $O(n \log n)$ time. Indeed, since a point and its nearest neighbor always define a Delaunay edge and the Delaunay triangulation has $O(n)$ edges, the nearest neighbor radii can be found in linear time. Once we know these radii, we can find the edges and faces in the sphere-of-influence diagram by performing a simple $O(1)$ test on the $O(n)$ edges and faces of the Delaunay triangulation.

The algorithm in Fig. 2 gives some details. The procedure *sid* extracts the sphere-of-influence diagram from the Delaunay triangulation of the point sample S and tags each extracted edge with the number of faces it belongs to: 2 for internal edges, 1 for external edges, and 0 for free edges. These tags are used to identify boundary points (see below); they may not be needed if a topological data structure is used. The procedure *nn-radii* (omitted) visits the edges of the Delaunay triangulation and fills a table R with the scaled nearest neighbor radius of each point: $R(p) = \mu r(p)$. These scaled radii are used in the predicate *sig* to decide whether an edge belongs to the sphere-of-influence graph (which is never fully built): two points define an edge when their scaled nearest neighbor disks intersect.

Extracting boundaries from the sphere-of-influence diagram. The triangulation of the sample points provided by the sphere-of-influence diagram is already useful for drawing the region and for deciding whether a query point in the plane is inside or outside it. Nevertheless, the boundaries of the sphere-of-influence diagram typically require much less data than the original sample points and thus provide a more concise description of the region. The procedure *boundaries* extracts the boundaries of a point sample S .

```

procedure sig( $u, v$ )
  return  $\text{dist}(u, v) \leq R(u) + R(v)$ 
end
procedure sid-edges( $E$ )
  for  $(u, v) \in E$  do
    if sig( $u, v$ ) then
      tag( $u, v$ )  $\leftarrow 0$ 
    else
      delete  $(u, v)$  from  $E$ 
    end
  end
end
procedure sid-faces( $F$ )
  for  $(u, v, w) \in F$  do
    if sig( $u, v$ ) and sig( $v, w$ ) and sig( $w, u$ ) then
      increment tag( $u, v$ ), tag( $v, w$ ), tag( $w, u$ )
    else
      delete  $(u, v, w)$  from  $F$ 
    end
  end
end
procedure sid( $S$ )
   $E, F \leftarrow \text{Delaunay}(S)$ 
  nn-radii( $E, \mu$ )
  sid-edges( $E$ )
  sid-faces( $F$ )
end
procedure boundaries( $S$ )
  sid( $S$ )
   $B \leftarrow \text{boundary-points}(E)$ 
  for  $v \in B$  do
    boundary-at( $v$ )
  end
end

```

Fig. 2. Extracting boundaries from the sphere-of-influence diagram.

After finding the sphere-of-influence diagram of S , it selects the boundary points as those adjacent to external edges. The boundary points are put into a set B . Then it finds the boundaries of the diagram by following the external edges of boundary points until it comes back to the starting point of each boundary polygon (code omitted). Boundary points are removed from B once they are visited to ensure that boundaries are traversed exactly once. This is a standard topological procedure.

Classifying external and internal boundaries. The procedure *boundaries* finds all boundaries of the diagram, but makes no distinction between external and internal boundaries. To classify a boundary polygon P as external or internal, consider its lowest vertex v (the one with smallest y -coordinate). If all vertices adjacent to v in the sphere-of-influence diagram are above v (that is, their y -coordinates are greater than or equal to that of v), then P is an external boundary; otherwise, P is an internal boundary. Once classified, the boundaries can be oriented consistently: counterclockwise for external boundaries, clockwise for internal boundaries. This orientation is important when filling polygonal




Fig. 3. Topological irregularities: free edges (top), split vertex (bottom); original region and samples (left), reconstructed region with $\mu = 1$ (middle) and with $\mu > 1$ (right): $\mu = 1.15$ (top), $\mu = 1.09$ (bottom).

regions using the non-zero winding rule, the default in the SVG image format. Considering P as a circular list, let u be the vertex in P before v and let w be the vertex in P after v . If P is an internal boundary, then reverse the list if $ccw(u, v, w)$. If P is an external boundary, then reverse the list unless $ccw(u, v, w)$. Here, $ccw(u, v, w)$ is the standard orientation predicate: it determines whether the triangle uvw is oriented counterclockwise by checking the sign of its oriented area, a 2×2 determinant.

Handling non-manifold boundary points. The procedure *boundaries* relies on each boundary point being adjacent to exactly two external edges. A boundary point that is adjacent to more than two external edges is a *non-manifold* boundary point (see Fig. 3). This topological irregularity in the sphere-of-influence diagram may occur when the sample is not well distributed. For topological robustness, we split each non-manifold boundary point into distinct topological vertices having the same position, one vertex for each fan of faces around the original non-manifold boundary point. This restores the requirement that each boundary point is adjacent to exactly two external edges and allows the algorithm to find the boundaries of all connected components, even if some components now touch. We flag the existence of non-manifold boundary points so that the user can tune μ to remove them.

Handling free edges. The other kind of topological irregularity that may occur when the sample is not well distributed are free edges joining parts of the diagram (see Fig. 3). By definition, free edges are not attached to faces and so cannot be boundary edges. Therefore, our algorithm never follows free edges. This may lead to extraneous connected components. The algorithm will find these components but not any isolated points. We flag the existence of free edges so that the user can tune μ to remove them if necessary.

Summary. To reconstruct a region in the plane from a point sample using our method, compute the sphere-of-influence diagram of the sample using $\mu = 1$ and extract the boundaries as described above. Our experience is that $\mu = 1$ suffices for spatially uniformly distributed sample. Otherwise, if the sample is less well distributed and topological irregularities are found

and are undesirable, increase $\mu > 1$ slowly until no topological irregularities exist. This is best done interactively or using binary search guided by the user. Note that the Delaunay triangulation and the table of unscaled nearest neighbor radii need to be computed only once. Moreover, the value of μ at which each Delaunay face appears in the sphere-of-influence diagram can be computed a priori. See the discussion on parameter tuning at the end of §5.

5. Experimental results and discussion

We now describe our experiments, starting with using $\mu = 1$ for reconstructing regions from spatially uniformly and smooth non-uniformly distributed samples (Figs. 4–9). Then we report how tuning $\mu > 1$ allows the reconstruction of regions from samples that are less well distributed, removing false holes and topological irregularities (Figs. 11–13). We also report qualitative and quantitative results that show that our method gives essentially the same results as existing methods: the widely used α -shapes [6] and the recent CT-shapes [7].

Tools and data. Here is the workflow used in our experiments. We computed Delaunay triangulations using Fortune’s algorithm and code [24] and also *qdelaulay* from the *qhull* suite [25]; both are command-line C programs. A program in Lua reads the points and the Delaunay triangles, computes the sphere-of-influence diagram (or alpha shape), and outputs a PostScript picture. We generated Poisson-disk point samples using the algorithm and code by Dunbar and Humphreys [26] and our own implementation of the algorithm by McCool and Fiume [27]. As data, we used country outlines extracted from the χ -shapes demo [28] and examples from the papers [7, 29].

Spatially uniformly distributed samples. For such samples, $\mu = 1$ works well. Fig. 4 shows an example reconstruction. Note how all connected components and holes are automatically identified and correctly reconstructed (in the sense that the topology is correct as perceived by a human). Fig. 5 shows additional examples. For a more extensive test, we generated samples inside polygonal regions describing country outlines by mapping a 7000-point blue noise sample in the square $[-1, 1] \times [-1, 1]$ linearly without distortion to a square containing the region, and then selecting the samples inside the region. Fig. 6 shows some examples.

Effect of sampling density. Fig. 7 illustrates how the reconstruction using $\mu = 1$ behaves when varying sampling densities. The original region is given implicitly by $f(x, y) = 0.004 + 0.110x - 0.177y - 0.174x^2 + 0.224xy - 0.303y^2 - 0.168x^3 + 0.327x^2y - 0.087xy^2 - 0.013y^3 + 0.235x^4 - 0.667x^3y + 0.745x^2y^2 - 0.029xy^3 + 0.072y^4 \leq 0$ [30]. We generated samples of different densities in the square $[-1, 1] \times [-1, 1]$ using Poisson-disk sampling for varying radii. Then we mapped the samples linearly without distortion to the square $[-1.6, 2.4] \times [-1.6, 2.4]$, which contains the region. Finally, we selected the samples inside the region by evaluating the implicit function. Fig. 7 shows the boundaries reconstructed from samples whose size is roughly halved each time. Fig. 8




Fig. 4. Using $\mu = 1$ works well for reconstructing regions from spatially uniformly distributed samples. All connected components and holes are automatically identified and correctly reconstructed.

shows that the reconstruction using $\mu = 1$ behaves well even with non-uniform densities, provided the density changes smoothly. In these examples, the sample points were extracted from the vertices of an adaptive triangular mesh. The example in Fig. 8 (right) differs from all other examples in that it contains both interior and boundary samples. The reconstruction works well because the boundaries and the interior are well sampled, and the interior samples are nicely away from the boundaries, respecting the density of the boundary samples. Fig. 9 shows that reconstruction using $\mu = 1$ may fail when the density changes abruptly; a correct reconstruction is found with $\mu > 1$.

Effect of sample quality. As illustrated in Fig. 7, sample density does not seem to affect the reconstruction using $\mu = 1$, as long as the sample is spatially uniform distributed: the reconstruction works equally well with samples of high and low density. On the other hand, the sphere-of-influence diagram is sensitive to the spatial distribution of points in the sample. When the points are not spatially uniformly distributed, small false holes may appear in the diagram when $\mu = 1$, even though they are not present in the hull (see Fig. 11). If no holes are expected in the region, then false holes can be ignored: the external boundaries of the components describe the region. Otherwise, we have to discriminate true holes from false ones. Closing holes with few vertices, specially three and four, or small area are natural choices, but we have not experimented with doing this automatically. Instead, we just tuned μ : increasing μ closes all false holes. Figs. 11–13 show examples where false holes appear when $\mu = 1$ and show the first $\mu > 1$ when false holes disappear and the correct topology of the region is found (again, correct in the sense of human perception). The coliseum example in Fig. 12 (bottom) uses stippling data [29], which has non-uniform but smoothly varying spatial distribution; again, false holes are closed and true holes are preserved. Fig. 13 also shows the first μ when the border is reasonably smooth visually. Higher values of μ may yield further smoothness, but if μ is too high, true holes will be closed and different connected components will be merged. (See the animation in the supplementary material.)

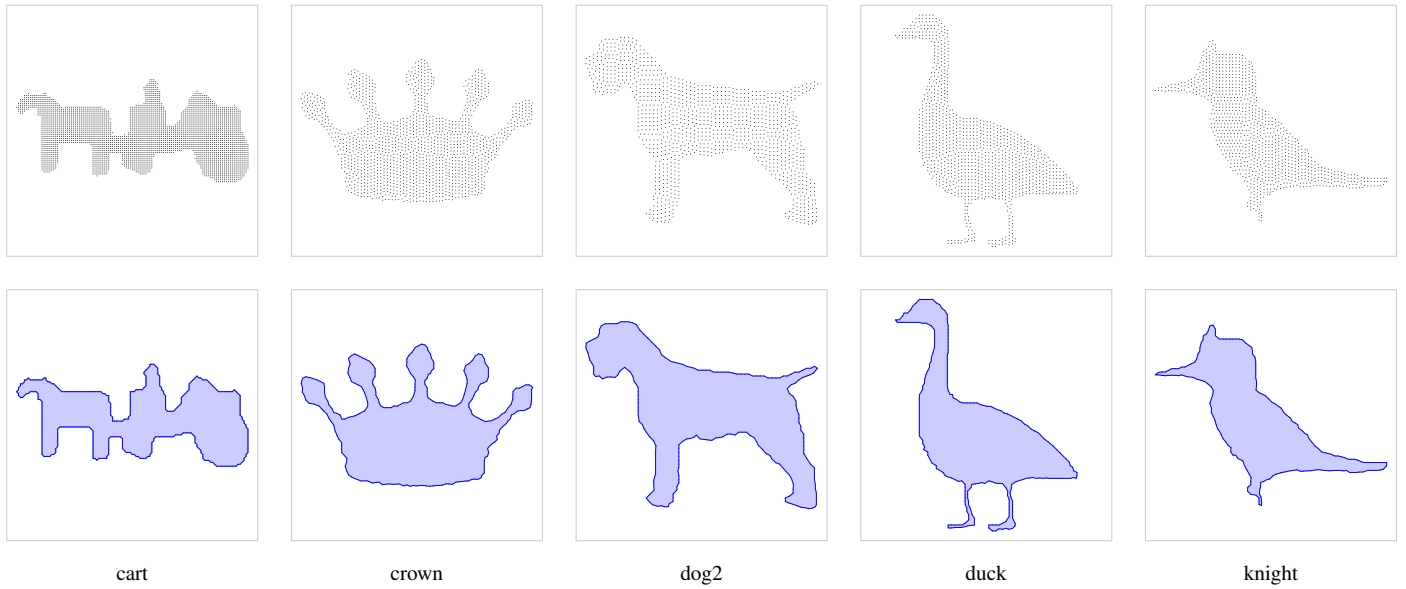


Fig. 5. $\mu = 1$ suffices for reconstructing regions from spatially uniformly distributed samples.

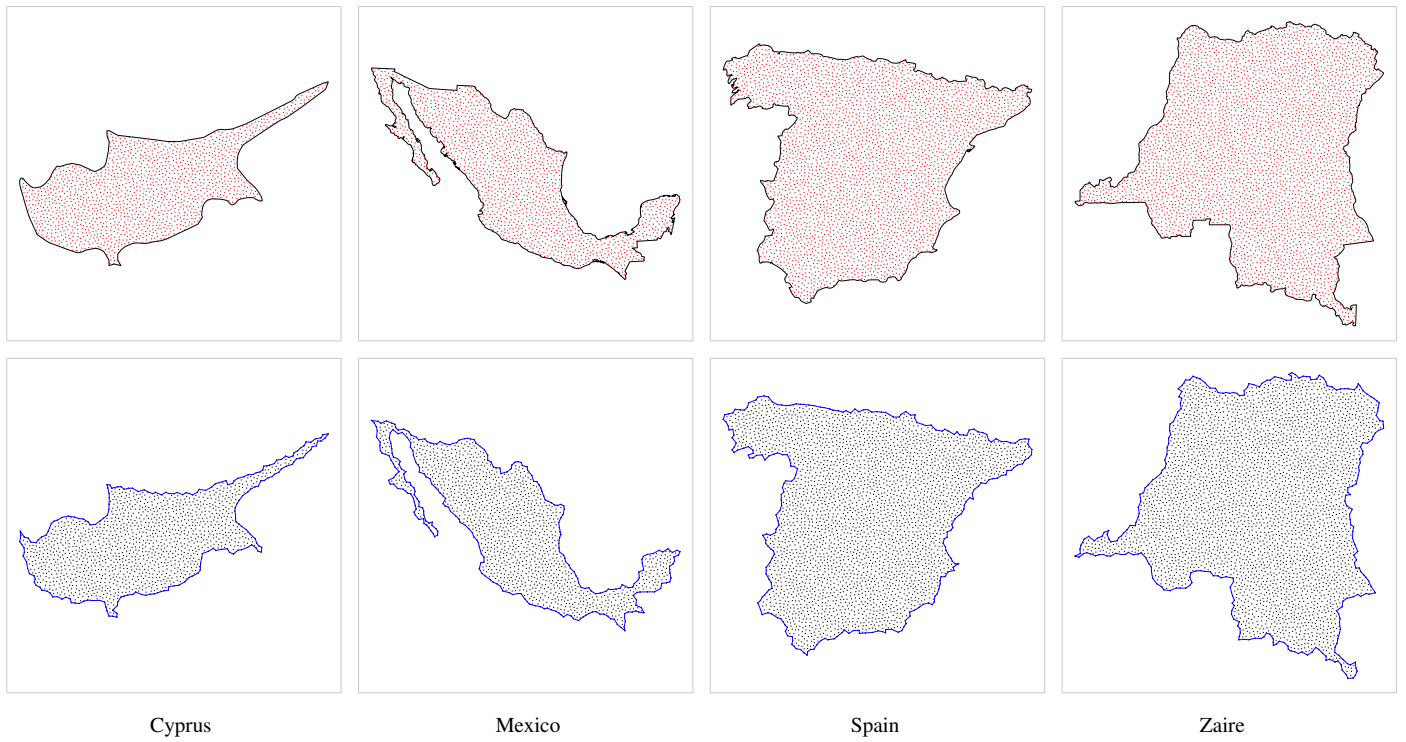


Fig. 6. Country outlines and samples (top); reconstructed regions using $\mu = 1$ (bottom).




Fig. 7. Effect of sampling density with $\mu = 1$. Poisson-disk sampling with $r = 0.005, 0.007, 0.010, 0.014, 0.020$ yielding $n = 7270, 3685, 1814, 944, 451$ points.




Fig. 8. Reconstruction behaves well with smooth non-uniform densities, even with $\mu = 1$.




Fig. 9. Reconstruction of non-smooth non-uniform densities: sample points (left), false holes with $\mu = 1$ (middle) do not appear when $\mu = 1.59$ (right).

Topological irregularities. Recall that two kinds of topological irregularities may occur when the sample is not well distributed: non-manifold boundary points and free edges (Fig. 3). We remove topological irregularities by increasing μ until they disappear (Figs. 11–13). We stress that, in our experiments reconstructing country outlines, topological irregularities are due to undersampling of the original geometry because the samples have fixed density, do not stay away from boundaries, do not include boundary samples, and so do not capture thin features.

Jagged boundaries. The boundary polygons found by our method are not smooth because their vertices are sample points that come from the interior of the region, not from its boundaries. Thus, the boundary vertices may be far from the actual (unknown) boundary. For country outlines, this is not too noticeable (Fig. 6). For regions with smooth boundaries, the reconstructed

boundaries may look jagged (Fig. 4), but can be improved using a method that extracts smooth curves from polygonal lines [31].

Qualitative comparison. Our results are visually the same as those obtained by other methods, such as α -shapes [6] and CT-shapes [7]. Fig. 14 shows some reconstructions using α -shapes; see also the supplementary material. They show the critical α and the first α and μ that recover the topology of the region. Although the reconstructions are quite similar, choosing an appropriate α is not automatic and there is no clear starting value. Choosing an appropriate μ is easier because we can start with $\mu = 1$, which frequently works already. See below for a discussion on parameter tuning.

The sphere-of-influence diagram tends to produce smoother boundaries, especially when $\mu > 1$. As a consequence, sometimes sample points that appear to a human as concave boundary vertices are smoothed out in the reconstruction. For instance, consider the intro example in Figs. 12 and 14. In the lower part where the tail joins the body, a natural concave boundary vertex appears in the data and in the reconstructions with $\alpha = 11.65$ and $\mu = 1$ but becomes an interior point when $\mu > 1$. This subtle behavior of the sphere-of-influence diagram may happen when the sample points are taken from a near-regular square grid. In this case, the edges of sphere-of-influence diagram typically join a sample point to its immediate grid neighbors in the sample, including the diagonal ones. Therefore, locally, the border of the sphere-of-influence diagram almost coincides with the orthogonal polygon containing the samples, except that the sphere-of-influence diagram prefers diagonal edges and so is slightly smoother (see also Fig. 10).

Quantitative comparison. We compared our results with those obtained by CT-shape for the point clouds shown in Figs. 5, 11–13. There is no ground truth for these point clouds, even though their shapes are easily perceived. We selected CT-shape as a proxy for ground truth because: it is a recent method that does not require parameter tuning; its code and data are readily available; and it has been extensively and favorably compared with previous methods. Table 1 shows that our reconstruction is quite close to the one by CT-shape. (Quite similar results would be obtained for α -shapes computed with a suitable value of α .) The *vertex* column shows the fraction of boundary vertices in

	μ	vertex	length	area
alien	1.26	0.987	0.968	1.004
cart	1.00	0.993	0.996	1.001
cathole	1.20	0.969	0.983	1.001
crown	1.00	0.996	0.958	1.012
deer	1.25	0.986	0.961	1.011
dog	1.46	0.991	0.956	1.014
dog2	1.00	0.991	0.975	1.007
dove	1.12	0.992	0.965	1.013
duck	1.00	0.996	0.973	1.008
fish	1.05	0.994	0.958	1.018
head	1.03	0.989	0.965	1.008
intro	1.42	1.000	0.943	1.007
jig	1.35	0.994	0.956	1.016
knight	1.00	0.995	0.974	1.008
multifish	1.31	0.997	0.950	1.014

Table 1. Quantitative comparison with CT-shape: fraction of coincident boundary vertices, ratio of boundary lengths, and ratio of region areas.

our reconstruction that appear in the reconstruction using CT-shape. The *length* column shows the ratio of the length of our boundary to the length of the boundary found by CT-shape. The *area* column shows the ratio of the area of our region to the area of the region found by CT-shape. We see that virtually all our vertices appear in CT-shape and the lengths and areas agree quite closely. This means that both boundaries are quite close to each other, except for minor local variations (see Fig. 10). Our boundaries are slightly shorter and our regions slightly larger because they tend to be smoother for $\mu > 1$.

For the country outlines, we had ground truth: the original polygons. However, the samples we generated have fixed density, are mostly away from the boundaries, and do not capture thin features. Thus, a direct comparison of boundaries would be meaningless. We opted to compare areas, with the expectation that the area of our reconstruction would be slightly less than the area of the original polygon. Nevertheless, the areas agree closely even when $\mu = 1$: the areas of all 154 reconstructed regions were at least 82% of the original. Moreover, the areas of 93% of all reconstructed regions were at least 90% of the original. For the regions in Fig. 6, the ratio of the areas are 0.92 for Cyprus, 0.90 for Mexico, and 0.96 for Spain and Zaire.

Parameter tuning. It is natural that a method for region reconstruction using Delaunay filtering includes a parameter t to control which edges and faces are removed. For concreteness, we follow α -shapes [6] and assume that, as we vary t from 0 to ∞ , the corresponding t -shapes grow monotonically from the isolated sample points to the full Delaunay triangulation. The t -spectrum of the sample is the finite ordered sequence of parameter values at which the shapes change topology. Focusing solely on the faces, we assign a *threshold* to each Delaunay triangle: the value of t at which the triangle first enters the shape. The spectrum is then the finite ordered sequence of thresholds. In α -shapes, the threshold of a Delaunay triangle is the radius of its circumcircle. In our method, the threshold of a Delaunay triangle uvw is $\mu(u, v, w) = \max(\mu(u, v), \mu(v, w), \mu(w, u))$, where

$$\mu(u, v) = \frac{\text{dist}(u, v)}{r(u) + r(v)}$$

and r is the unscaled nearest neighbor radius.




Fig. 10. Comparing our reconstruction using $\mu = 1.03$ (blue) with CT-shape (red). Both boundaries follow each other quite closely, except for minor local variations. (Part of the head example shown in Fig. 12.)

The key task is how to choose a suitable value of t to reconstruct the region. This is typically done interactively or using binary search guided by the user. To help in this task, we can analyze the spectrum of the sample and take its *critical value*: the smallest parameter value for which the shape includes all sample points (a notion borrowed from the α -shape toolbox in MATLAB). The critical value is a good starting point for region reconstruction, but frequently cannot be used directly because including all points does not ensure the absence of false holes.

Although α -shapes are widely used for shape reconstruction, finding a suitable α is challenging because the α -spectrum typically has values of widely different scales, including very large values corresponding to thin triangles near the boundary of the convex hull. (For the point sample in Fig. 4, the α -thresholds range from 5.73 to 52923.99; the critical value is 9.85.) Moreover, the α -spectrum is not scale-invariant. In contrast, the μ -spectrum is scale-invariant, which accounts for its much reduced range: the largest μ is essentially the length of the longest Delaunay edge, normalized by the distance between neighboring points. (For the point sample in Fig. 4, the μ -thresholds range from 0.50 to 20.80; the critical value is 0.94.) Moreover, our method starts with $\mu = 1$, which already gives a fairly good reconstruction, frequently a topologically correct reconstruction for well-distributed samples. When there are false holes, a suitable μ is not very far from $\mu = 1$ (see Figs. 11–13).

6. Conclusion

The sphere-of-influence diagram is a natural subset of the Delaunay triangulation that works remarkably well for reconstructing regions from spatially uniformly distributed point samples, already for $\mu = 1$. It automatically identifies connected components and holes, which seems to be the hardest algorithmic task but the easiest perceptual one. When the sample has lower spatial quality, tuning the scaling parameter μ helps to close

false holes and to handle topological irregularities. Tuning μ also improves the smoothness of the reconstructed boundaries. Choosing a suitable μ for the sphere-of-influence diagram is easier than choosing a suitable α for α -shapes because we can start at $\mu = 1$ and the μ -spectrum has a much reduced range. (The μ -spectrum is scale-invariant, whereas the α -spectrum is not.) Like α -shapes, exploring the μ -spectrum interactively is efficient because the Delaunay triangulation and the μ -thresholds need to be computed only once. Finally, the sphere-of-influence diagram is easily extracted in linear time from the Delaunay triangulation: it relies on a simple geometric predicate that removes Delaunay edges and faces. The whole process takes thus $O(n \log n)$ asymptotic time. This coincides with the complexity of other methods, which use at most $O(n \log n)$ time to perform Delaunay filtering. We expect that these qualities will make our region reconstruction method attractive in applications, such as geometric processing of geographic queries and pattern recognition from images.

Further work. A natural direction for further work is to relax the requirements on the output. The description of the boundaries by polygonal lines whose vertices come only from the point sample is not a hard requirement for reconstructing the region. Relaxing this requirement would allow boundaries described by simpler polygonal lines and nice smooth splines. The requirement that every sample point is inside exactly one contour is a harder requirement. Relaxing this requirement depends on defining and identifying outliers in the sample and may be useful for handling noisy samples.

Regarding the input, our experiments show that the sphere-of-influence diagram reconstructs the boundaries of a region from a well-distributed sample. We use ‘well-distributed’ informally. A theoretical analysis that made this hypothesis precise and provided guarantees that our method works for some concrete classes of sample distributions, such as Poisson-disk sampling, would nicely complement our experimental work. A natural starting point is to consider r -samplings [21, 22], perhaps amended to ensure good separation of interior samples.

Acknowledgements. We thank the authors of the papers [19, 7, 29] for generously sharing their data with us. Afonso Paiva is partially supported by CNPq under grant 309226/2020-1 and by FAPESP under grant 2019/23215-9.

References

- [1] Edelsbrunner, H. Shape reconstruction with Delaunay complex. In: LATIN’98; vol. 1380 of *Lecture Notes in Computer Science*. 1998, p. 119–132.
- [2] Asaedi, S, Didehvar, F, Mohades, A. α -concave hull, a generalization of convex hull. *Theoretical Computer Science* 2017;702:48–59.
- [3] Toussaint, GT, editor. Computational Morphology; vol. 6 of *Machine Intelligence and Pattern Recognition*. North-Holland; 1988.
- [4] Toussaint, GT. A graph-theoretical primal sketch. In: [3]. 1988, p. 229–260.
- [5] Toussaint, GT. The sphere of influence graph: theory and applications. *International Journal of Information Technology & Computer Science* 2014;14(2):37–42.
- [6] Edelsbrunner, H, Kirkpatrick, D, Seidel, R. On the shape of a set of points in the plane. *IEEE Transactions on Information Theory* 1983;29(4):551–559.
- [7] Thayyil, SB, Parakkat, AD, Muthuganapathy, R. An input-independent single pass algorithm for reconstruction from dot patterns and boundary samples. *Computer Aided Geometric Design* 2020;80:101879. Code available at github.com/agcl-mr/Reconstruction-CTShape.
- [8] Kirkpatrick, DG, Radke, JD. A framework for computational morphology. In: Toussaint, GT, editor. *Computational Geometry*; vol. 2 of *Machine Intelligence and Pattern Recognition*. North-Holland; 1985, p. 217–248.
- [9] Radke, JD. On the shape of a set of points. In: [3]. 1988, p. 105–136.
- [10] Medeiros, E, Velho, L, Lopes, H. Restricted BPA: applying ball-pivoting on the plane. In: *SIBGRAPI 2004*. IEEE Press; 2004, p. 372–379.
- [11] Klein, J, Zachmann, G. Point cloud surfaces using geometric proximity graphs. *Computers & Graphics* 2004;28(6):839–850.
- [12] Vital Brazil, EA, de Figueiredo, LH. A heuristic method for region reconstruction from noisy samples. *International Journal of Shape Modeling* 2009;15(1-2):1–17.
- [13] Galton, A, Duckham, M. What is the region occupied by a set of points? In: *Geographic Information Science*; vol. 4197 of *Lecture Notes in Computer Science*. Springer; 2006, p. 81–98.
- [14] Jarvis, R. On the identification of the convex hull of a finite set of points in the plane. *Information Processing Letters* 1973;2(1):18–21.
- [15] Jarvis, R. Computing the shape hull of points in the plane. In: *Proceedings of IEEE Computer Society Conference on Pattern Recognition and Image Processing*. IEEE; 1977, p. 231–241.
- [16] Duckham, M, Kulik, L, Worboys, M, Galton, A. Efficient generation of simple polygons for characterizing the shape of a set of points in the plane. *Pattern Recognition* 2008;41(10):3224–3236.
- [17] Zhong, X, Duckham, M. Characterizing the shapes of noisy, non-uniform, and disconnected point clusters in the plane. *Computers, Environment and Urban Systems* 2016;57:48–58.
- [18] Zhong, X, Duckham, M. An efficient incremental algorithm for generating the characteristic shape of a dynamic set of points in the plane. *International Journal of Geographical Information Science* 2017;31(3):569–590.
- [19] Peethambaran, J, Muthuganapathy, R. A non-parametric approach to shape reconstruction from planar point sets through Delaunay filtering. *Computer-Aided Design* 2015;62:164–175.
- [20] Methirumangalath, S, Parakkat, AD, Muthuganapathy, R. A unified approach towards reconstruction of a planar point set. *Computers & Graphics* 2015;51:90–97.
- [21] Methirumangalath, S, Kannan, SS, Dev Parakkat, A, Muthuganapathy, R. Hole detection in a planar point set: An empty disk approach. *Computers & Graphics* 2017;66:124–134.
- [22] Thayyil, SB, Peethambaran, J, Muthuganapathy, R. A sampling type discernment approach towards reconstruction of a point set in R^2 . *Computer Aided Geometric Design* 2021;84:101953.
- [23] Imai, H, Iri, M, Murota, K. Voronoi diagram in the Laguerre geometry and its applications. *SIAM Journal on Computing* 1985;14(1):93–105.
- [24] Fortune, S. A sweepline algorithm for Voronoi diagrams. *Algorithmica* 1987;2(1-4):153–174. Code available at 9p.io/who/sjf/.
- [25] Barber, CB, Dobkin, DP, Huhdanpaa, H. The quickhull algorithm for convex hulls. *ACM Transactions on Mathematical Software* 1996;22(4):469–483.
- [26] Dunbar, D, Humphreys, G. A spatial data structure for fast Poisson-disk sample generation. In: *ACM SIGGRAPH 2006*. 2006, p. 503–508. Code available at github.com/ddunbar/PDSample.
- [27] McCool, M, Fiume, E. Hierarchical Poisson disk sampling distributions. In: *Graphics Interface ’92*. 1992, p. 94–105.
- [28] Hudson, G, Duckham, M. Characteristics shapes. 2008. URL: <http://duckham.org/matt/characteristics-shapes/>.
- [29] Kaplan, CS, Bosch, R. TSP art. In: *Renaissance Banff: Mathematics, Music, Art, Culture*. Bridges Conference; 2005, p. 301–308.
- [30] Taubin, G. Distance approximations for rasterizing implicit curves. *ACM Transactions on Graphics* 1994;13(1):3–42.
- [31] Plass, M, Stone, M. Curve-fitting with piecewise parametric cubics. *ACM SIGGRAPH Computer Graphics* 1983;17(3):229–239.



Fig. 11. First μ that recovers the correct topology of the region (last column).

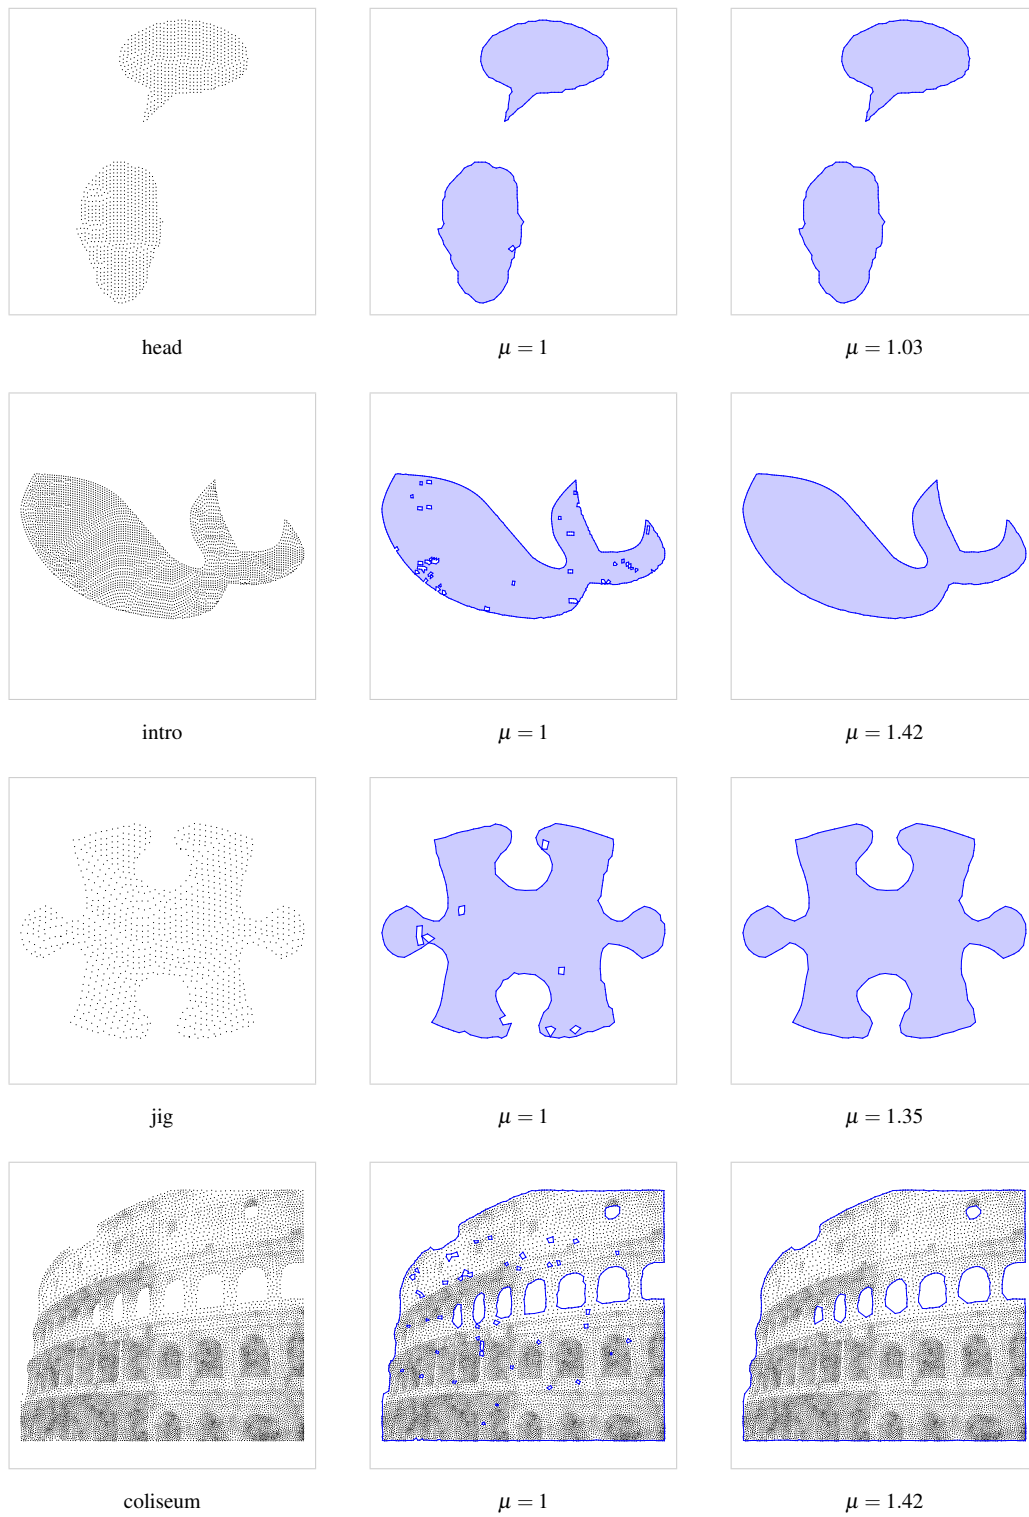


Fig. 12. First μ that recovers the correct topology of the region (last column).



Fig. 13. First μ that recovers topology (third column); next μ that gives smooth reconstruction (fourth column).

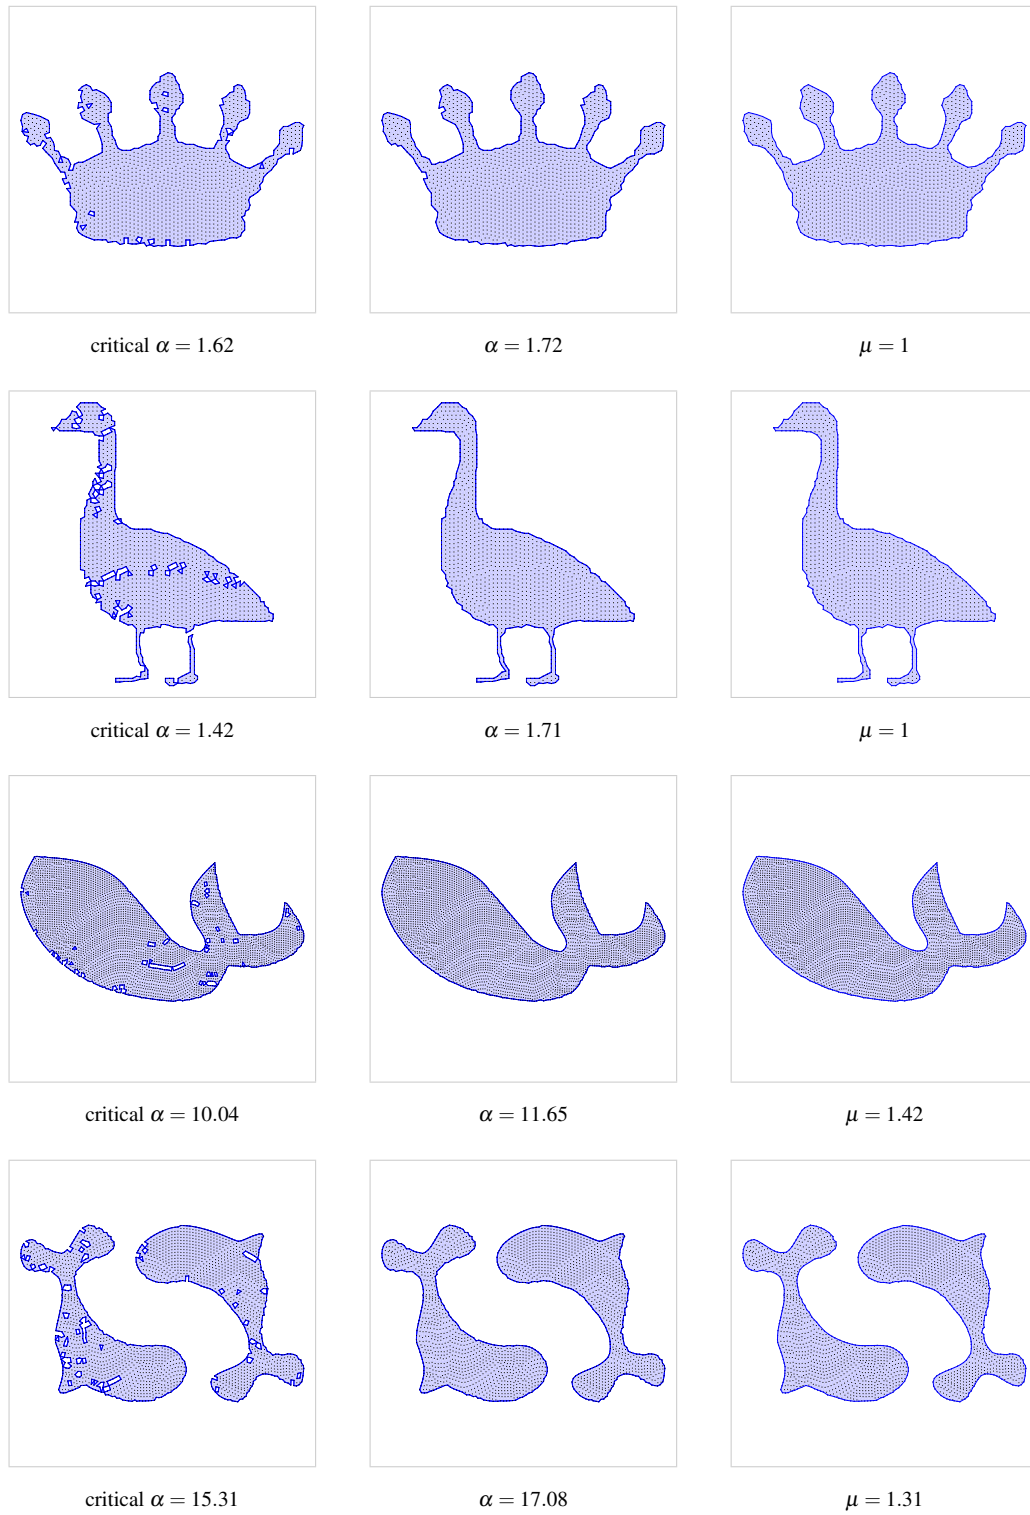


Fig. 14. Critical α and the first α and μ that recover the topology of the region.