Regular and non-regular point sets: properties and reconstruction

Sylvain Petitjean, Edmond Boyer

Regular and Non-Regular Point Sets: Properties and Reconstruction

Sylvain Petitjean [1]

Loria–CNRS & Inria Lorraine, Campus scientifique, BP 239, 54506 Vandœuvre-les-Nancy cedex, France

Edmond Boyer [2]

Gravir–Inria Rhône-Alpes, 655 Avenue de l'Europe, 38330 Montbonnot, France

Abstract

In this paper, we address the problem of curve and surface reconstruction from sets of points. We introduce regular interpolants, which are polygonal approximations of curves and surfaces satisfying a new regularity condition. This new condition, which is an extension of the popular notion of r-sampling to the practical case of discrete shapes, seems much more realistic than previously proposed conditions based on properties of the underlying continuous shapes. Indeed, contrary to previous sampling criteria, our regularity condition can be checked on the basis of the samples alone and can be turned into a provably correct curve and surface reconstruction algorithm. Our reconstruction methods can also be applied to non-regular and unorganized point sets, revealing a larger part of the inner structure of such point sets than past approaches. Several real-size reconstruction examples validate the new method.

Key words: Surface reconstruction, sampling condition, regular interpolants

1 Introduction

Computing a polygonal approximation of a curve or a surface from a set of unorganized sample points is a challenging problem in several fields such as computer vision, computer graphics and computational geometry. The concerned applications are, for instance, curve reconstruction in image analysis,

[1] E-mail: Sylvain.Petitjean@loria.fr
[2] E-mail: Edmond.Boyer@inrialpes.fr

three-dimensional model determination from laser range data and from computer vision processes (stereo-vision, reconstruction from apparent contours, ...), reconstruction in medical imaging and the creation of computer models from existing parts in reverse engineering.

The problem faced in these applications can be stated in the following way: given a set of points of a plane curve or of a three-dimensional surface, construct a polygonal structure interpolating the sample points that reasonably captures the shape of the point set. In this paper, we introduce a new, simple and efficient algorithm for reconstructing curves and surfaces from unorganized sets of points.

1.1 Past work

Shape reconstruction from sets of points has been widely studied in the computer graphics and computational geometry communities in recent years. We here give an overview of the existing reconstruction methods, using the classification established by Mencl and Muller [1].

Spatial subdivision. Common to all spatial subdivision techniques is the subdivision of a bounding volume surrounding the input data into disjoint cells. The goal of these algorithms is to combinatorially select cells related to the shape of the point set. Most algorithms in this category are based on the Delaunay triangulation (2D) or tetrahedrization (3D) of the input data. An early reconstruction algorithm of this kind is the Delaunay "sculpting" heuristic of Boissonnat [2]: tetrahedra are progressively eliminated according to a geometric criterion taking into account the areas of their faces. In [3], Veltkamp computes γ-graphs in a similar way, though he eliminates tetrahedra using a slightly different criterion which has the advantage of adapting to variable point density. Both of these approaches cannot handle objects with holes and surface boundaries. The methods based on the well known α-shapes fall also in this category. Recall that α-shapes, introduced by Edelsbrunner and Mücke [4], represent a finite set of points at different levels of detail. The main difficulty appears that the best reconstruction may require different α's in different places. Recently, Guo et al. [5] have tried to improve the results of reconstruction with α-shapes using visibility algorithms. In the 2D case, the method proposed by Attali [6], based on the notion of *normalized meshes*, overcomes some of the drawbacks of the α-shape algorithms.

Also in the category of volume-based cell selection, Amenta et al. [7,8] use a Voronoi filtering approach based on 3D Voronoi diagrams and Delaunay triangulations, to construct an interpolating shape of the sample points called the *crust*. In [9], a simplified version of the reconstruction algorithm is proposed.

Zero-set approximation. In the zero-set type of approaches, the input points are used to define a signed distance function on $\mathbb{R}^3$, the zero-set of which is then polygonalized to create the output mesh. Compared to spatial subdivision techniques, this approach approximates rather than interpolates sample points. Hoppe et al. [10] determine an approximate tangent plane at each sample point and then use the distance to the nearest point's tangent plane as a distance function to the surface. The zero-set of this function is then computed by the Marching Cubes algorithm, which outputs a mesh in triangular form. Curless and Levoy [11] use a similar algorithm, but for laser range data, from which they derive tangent plane information. The input samples are combined into a volumetric function which is computed and stored on a voxel grid. These algorithms require a uniform sampling, at least locally. Bernardini and Bajaj [12] propose an algorithm that computes a distance function from the α-shape of the points. The α-shape has the same topology as the actual surface only if the sampling is dense and uniform.

A promising new work in this category, which combines Voronoi diagrams and implicit functions, is the natural neighbor interpolation approach of Boissonnat and Cazals [13]. The authors present an algorithm which, starting with a set of unorganized points equipped with normal directions, reconstructs a smooth surface of arbitrary topology. The surface is implicitly represented as the zero-set of a signed pseudo-distance function. The interpolation of non-uniform samples is constructed using the notion of natural neighbors, which can be computed efficiently from the Voronoi diagram of the sample. Theoretical guarantees are derived for the quality of the reconstructed surface.

Surface deformation. Reconstruction methods in this category deform an initial surface to give a good approximation of the input point set. For instance, the desired shape can be considered an elastic membrane. The starting point is a large membrane enclosing the input points. A deformation process is then applied to minimize the energy until a local minimum is reached. If the initial guess is not too far from the desired shape, these methods are fast and robust against noise. The variational level set formulation of Osher et al. [14] has proved very effective in this respect.

Incremental surface construction. The idea behind incremental construction schemes is to build up the interpolating mesh using surface-oriented properties of the input sample points. This can be done, as in [2], by selecting an initial edge connecting two points expected to be neighbors on the surface and then iteratively extending the edge to a large surface by attaching further triangles at the boundaries of the constructed surface. Surface construction can also start with a global wireframe of the surface which is iteratively filled to complete the surface. Mencl and Müller [15] use the Euclidean minimum spanning tree as initial wireframe.

Recent advancing-front triangulation schemes include the Ball-Pivoting Algorithm (BPA) of Bernardini et al. [16], the Spiraling-Edge triangulating technique of Crossno and Angel [17] and the projection-based method of Gopi and Krishnan [18]. Given a point set and a radius r, the BPA finds an interpolating mesh where each of the triangles is characterized by the fact that the ball of radius r that sits on its vertices has no internal point. The algorithm works by finding a seed triangle, then extending the surface as far as it can by pivoting a ball of radius r along each boundary edge of the current surface. Under some sampling conditions, the BPA is guaranteed to finish with a correct triangulation. The Spiraling-Edge technique works by creating a star-shaped triangulation between a point and its neighbors. In the method of Gopi and Krishnan, a local Delaunay neighborhood is computed around each point of a set of candidate points and the final surface triangulation is determined from this neighborhood relationship. These algorithms are closely related. But the latter one appears capable of handling very large data sets in a small amount of time. Also, a topologically correct mesh is guaranteed after surface reconstruction.

1.2 This paper

Dealing with surfaces of arbitrary topology, allowing non-uniform sampling and producing models with provable guarantees are currently the most crucial issues in surface reconstruction. Assuming that the input data points are sampled from some actual surface, a correct reconstruction of this surface is possible only if it is "properly" sampled, i.e. if there are enough sample points in areas of high curvature variation. But it is also desirable to avoid oversampling in low curvature regions.

Recently, emphasis has been put on the definition of "good" samplings of a shape and on the development of algorithms with provable guarantees, i.e. such that the reconstruction is guaranteed to be topologically correct and convergent to the original surface as the sampling density increases. Attali [6] and Faugeras [19] define a sampling condition in the 2D case which is global over the shape to reconstruct and thus does not account for local shape variations. Bernardini and Bajaj [20] also give sufficient conditions on the sampling to allow a homeomorphic, error-bounded reconstruction. The algorithm of Amenta et al. [8], based on the notion of r-sampling, is the first one with provable guarantees in 3D. The definition of a "good" sampling used is based on a local analysis and captures the intuitive notion that featureless areas can be reconstructed from fewer samples. A similar criterion is used by Boissonnat and Cazals [13] to give theoretical guarantees for natural neighbor interpolation. Gopi and Krishnan [18] present a new local sampling condition, based on estimations of directional curvatures at sample points.

These theoretical results hold when the sampling of the underlying surface is "sufficiently" dense. However, such sampling conditions are rarely met in practical applications. One main reason is that sampling criteria put too much emphasis on the original curve or surface from which the sample points are drawn and not enough on properties of the the interpolating mesh. However, the fact that connections between adjacent points can be recovered depends much more on local properties of the reconstructed surface than on properties of the original shape.

This motivated our definition of a discrete equivalent of r-sampling (based on an analogue for piecewise-linear objects of the medial axis used in the continuous case) and our study of the properties of the interpolating meshes satisfying this new regularity condition, called *regular interpolants*. Interestingly, the regularity of a point set (i.e., the fact that it admits a regular interpolant) can be decided on the basis of the points alone, by contrast to the original r-sampling criterion. In turn, this translates into a powerful algorithm for reconstructing the interpolant of a regular point set (Proposition 12 and Algorithm 2 in the 3D case).

The new method can also be applied to non-regular and unorganized point sets. The idea then is to iteratively determine locally regular configurations of simplices. In that case, more simplices will be computed than is really needed, revealing the internal structure of the data, and a further step based on heuristics is required to extract interesting parts from this structure. Experiments prove that our algorithms recover more desired connections than previous reconstruction approaches based on an elimination strategy. We have applied our reconstruction method to different types of data, in particular real data coming from computer-vision reconstruction process, and the results obtained show the correctness of the approach.

The paper is organized as follows. Section 2 introduces definitions and geometric structures that are used in the rest of the paper. In particular, a key tool of our work is an extension to polyhedral shapes of the notion of medial axis. Fundamental properties of regular interpolants are studied in Section 3. Regular interpolants of point sets of $\mathbb{R}^2$ are examined in Section 4 and a reconstruction method is proposed. Section 5 then moves up to the 3D case. And Section 6 shows that the properties of regular interpolants can be turned into heuristics for reconstructing non-regular point sets. Implementation details and significant results of reconstruction are given in Section 7, before concluding.

2 Definitions

If a smooth plane curve F is sampled densely enough, a graph on the set of sample points $\mathcal{P}$ can be constructed which is topologically consistent with the curve. This is the result that Amenta et al. [7,21] originally proved when F is r-sampled, i.e. when the distance of each point $P \in F$ to $\mathcal{P}$ is smaller than the local feature size at P (distance between P and the medial axis of F) multiplied by r, where r is something like .2 in 2D and .05 or even smaller in 3D.

Unfortunately, this result (and related sampling theorems) does not extend to practical situations where only sample points are given. Accordingly, our aim in this section is to introduce an extension of the notion of r-sampling to the discrete case. Sections 3 to 6 will then show what this discrete definition leads to in terms of reconstruction.

The rest of this section is organized as follows. We start (§ 2.1) by defining the notion of interpolant of a point set. We then introduce analogues of the distance to the sample point set (§ 2.2) and of the local feature size (§ 2.4), which we call respectively the local granularity and the local thickness. The latter is based on the notion of discrete medial axis of an interpolant (§ 2.3). Our definition of regularity of an interpolant, which is an extension of r-sampling to the discrete case, follows (§ 2.5).

In what follows, let $\mathcal{P}$ be a finite set of points of $\mathbb{R}^m$, $m = 2$ or 3.

2.1 Interpolant

Recall that a point, an edge, a triangular region and a tetrahedral region are respectively called a 0-simplex, a 1-simplex, a 2-simplex and a 3-simplex. A *k-simplicial complex* K of $\mathbb{R}^m$ is a collection of l-simplices, $l \in \{0, \ldots, k\}$ such that every face of a simplex of K is in K and such that the intersection of any two simplices of K is a face of each of them.

We first define what we mean by a piecewise-linear hypersurface of $\mathbb{R}^m$ (curve in 2D, surface in 3D) interpolating the points of $\mathcal{P}$:

Definition 1 (Interpolant) *A* piecewise-linear interpolant $\mathcal{O}$ *of* $\mathcal{P}$ *(or simply* interpolant *for short) is a pure* $(m-1)$*-simplicial complex whose vertex set is* $\mathcal{P}$*. The interpolant is said to be* closed *if each of its* $(m-1)$*-simplices borders two different connected components of* $\mathbb{R}^m \setminus \mathcal{O}$*.*

This definition is illustrated by Figure 1. Note that dangling edges are not

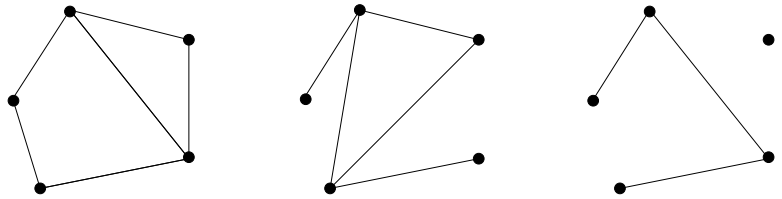

Fig. 1. The left complex is a closed interpolant, the center one an open interpolant and the right one is not considered an interpolant.

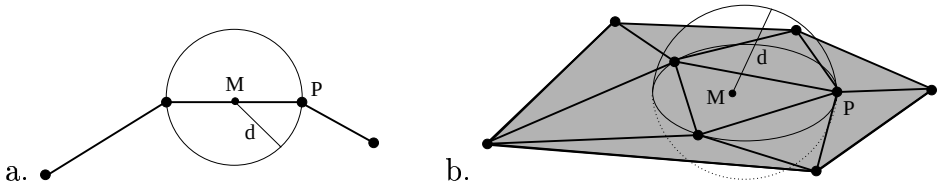

Fig. 2. The local granularity of an interpolant at vertex P. a. For a plane curve. b. For a surface.

allowed in closed interpolants.

2.2 Local granularity

From now on, assume that $\mathcal{O}$ is a closed interpolant of $\mathcal{P}$. Recall that the *star* of a vertex P of a simplicial complex is the set of simplices incident on P. The local granularity, as the distance to the sample point set in the continuous case, is a local measure of the density of the sampling.

Definition 2 (Local granularity) *The* local granularity $g_{\mathcal{O}}(P)$ *of* $\mathcal{O}$ *at* P *is the radius of the largest of the balls circumscribed to the simplices in the star of* P*.*

Examples of local granularities are shown in Figure 2.

2.3 Discrete medial axis

The very intuitive notion of medial axis, also called skeleton in the literature, can be defined in several ways which do not always lead to the same object: only regularity conditions on the shape will blur the differences.

Here we use the definition by maximal open balls. Let X be an open set of $\mathbb{R}^m$ with a smooth boundary. The *medial axis* of X is the closure of the locus of centers of open balls that are maximal with respect to inclusion in X [22, p. 376]. Also, the medial axis of the boundary ∂X of X may be defined as the union of the medial axes of X and of the complement of its closure $\overline{X}^c$.

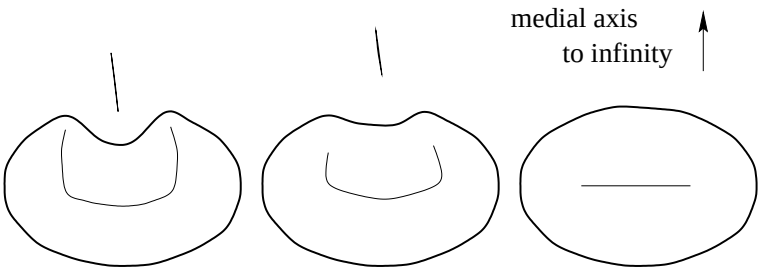

Fig. 3. The convexification of a shape sends its outside medial axis to infinity.

The medial axis of ∂X so defined has several important properties:

- each connected component of $\mathbb{R}^m \setminus \partial X$ contains a piece of the medial axis (we may assume that if X is convex, $\overline{X}^c$ has its medial axis at infinity – see Fig. 3) and this piece is connected;
- let P be any point of ∂X and C a component of $\mathbb{R}^m \setminus \partial X$ whose closure contains P; then P "contributes" to the part of the medial axis of ∂X due to C, i.e. it lies on the boundary of at least one maximal ball of C;
- the medial axis of ∂X nowhere intersects ∂X.

To define an analogue of r-sampling in the discrete case, we need to introduce a definition of the medial axis of an interpolant (hereafter referred to as a discrete medial axis) that closely parallels the above definition in the continuous case and has similar properties. Several methods have been proposed for defining the medial axis of a piecewise-linear surface (see, e.g., [23–26]). Most are based on the Voronoi diagram of the underlying vertex set. None, however, is entirely satisfactory for the present purpose.

We thus introduce a new definition of medial axis (which is related to the one proposed in [23]). Call a maximal open ball of $\mathbb{R}^m \setminus \mathcal{P}$ a Voronoi ball of $\mathcal{P}$. If B is a Voronoi ball, denote by $\rho(B)$ the set of samples of $\mathcal{P}$ lying on the boundary of B. Call B *free* if $\rho(B)$ is not the vertex set of a simplex of $\mathcal{O}$, i.e. if the samples in $\rho(B)$ are not cosimplicial.

Definition 3 (Discrete medial axis) *The* medial axis *of $\mathcal{O}$ is the closure of the locus of centers of free Voronoi balls.*

In other words, the medial axis of $\mathcal{O}$ is the Voronoi diagram of $\mathcal{P}$ minus those parts contributed by cosimplicial points of $\mathcal{O}$. With this definition, we see that the medial axis of a triangle in 2D is not empty but is reduced to a single point: the circumcenter of the triangle. The same goes in 3D. On the right of Fig. 4, P, Q and R are pairwise 1-cosimplicial but not 2-cosimplicial (i.e. vertices of the same 2-simplex), so any maximal ball through P, Q and R is free.

The discrete medial axis so defined has several of the interesting properties

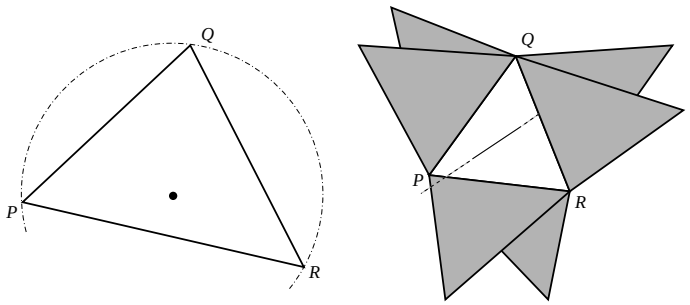

Fig. 4. Examples of discrete medial axes (thin lines) of interpolants (thick lines). Left (2D): the medial axis is reduced to a single point, the circumcenter of the triangle PQR. Right (3D): part of the medial axis is located on the intersection of the three median planes of the triangle PQR.

of the continuous medial axis, but obviously not all. Consider the example of Fig. 5.a: because the sample point set is not "dense enough" to capture all the details of the curve, the medial axis of the inside of the interpolant has been pushed outside of the shape by going discrete. Another example of bad sampling is shown on Fig. 5.b: because of the poor configuration of vertices, the inside medial axis crosses the interpolant and goes to infinity.

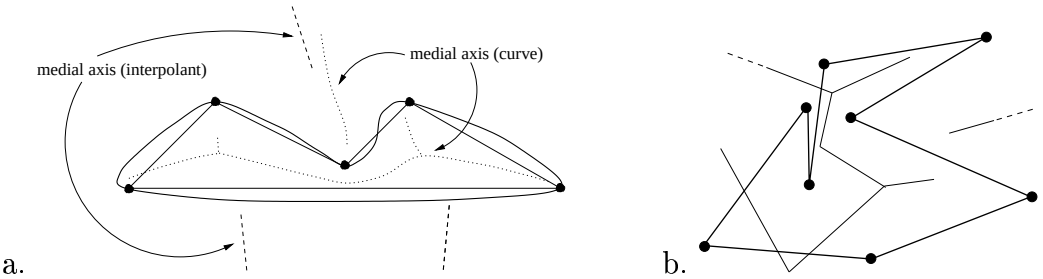

Fig. 5. Two examples of discrete medial axes and how they differ from the continuous case.

What these two examples tend to indicate is that the discrete medial axis is a good indicator of the ability of an interpolant to reconstruct an unknown shape. By imposing a condition on the distance of a vertex to the medial axis, we will build a class of interpolants (called regular interpolants) whose medial axis is very much like the continuous medial axis.

2.4 Local thickness

To convey the above-mentioned properties of the open set medial axis to the discrete case, we need additional hypotheses on interpolants. These hypotheses are based on the notion of local thickness, which we introduce now.

Definition 4 (Local thickness) *Let $P \in \mathcal{P}$ be a vertex of $\mathcal{O}$ and $\mathcal{C}_1, \ldots, \mathcal{C}_k$ the components of $\mathbb{R}^m \setminus \mathcal{O}$ whose closures contain P. Consider the contribution*

of P to the medial axis of $\mathcal{O}$ and let d_i be the distance from P to the part of this contribution that falls inside $\mathcal{C}_i$ (d_i may be ∞ if the component is unbounded). If d_i is undefined (no contribution of P to the medial axis inside $\mathcal{C}_i$), set $d_i = 0$.

The local thickness $e_{\mathcal{O}}(P)$ *of $\mathcal{O}$ at P is defined as:*

$$e_{\mathcal{O}}(P) = \min_{i=1,\ldots,k} d_i.$$

Two examples of local thicknesses are shown on Fig. 6, using the notations of Definition 4. Note that this definition is similar to the local feature size.

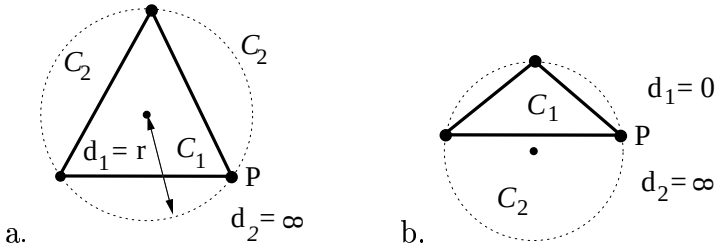

Fig. 6. Local thickness examples. $\mathcal{C}_1$ and $\mathcal{C}_2$ are the components of $\mathbb{R}^m \setminus \mathcal{O}$. The local thickness at P is: a. $e_{\mathcal{O}}(P) = r$. b. $e_{\mathcal{O}}(P) = 0$.

2.5 Regular interpolant

Now that granularity and thickness have been introduced, we can define what we mean by a "good" interpolant.

Definition 5 (Regular interpolant) *$\mathcal{O}$ is said to be a* regular interpolant *of $\mathcal{P}$ if at each point P of $\mathcal{P}$ the local granularity is strictly smaller than the local thickness:*

$$g_{\mathcal{O}}(P) < e_{\mathcal{O}}(P).$$

The point set $\mathcal{P}$ is called regular *if it admits at least one regular interpolant. A regular interpolant of a regular point set is called a* realization *of the point set.*

A few examples of regular and non-regular point sets are shown in Fig. 7.

This definition of regularity is also similar to the definition of a r-sampled curve given by Amenta et al. The major difference, of course, is that the r-sampling condition, which gives a criterion for a curve to be reconstructed from a set of samples, makes explicit reference to the medial axis of the curve which, given only samples, is not known and cannot be inferred. In other words, the r-sampling notion says nothing about what can be recovered from

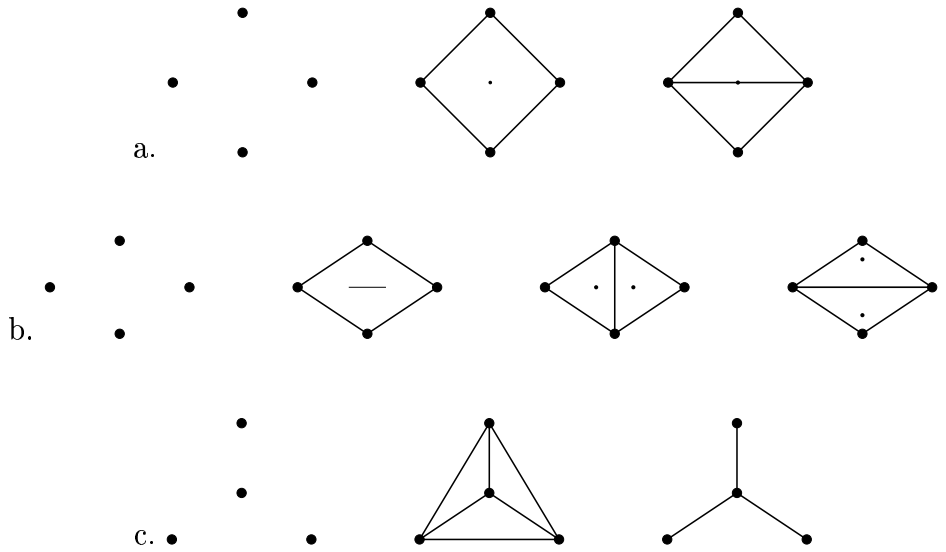

Fig. 7. Examples of regular and non-regular point sets. a. A regular point set, its sole realization and a non-regular interpolant. b. A regular point set, its two realizations and a non-regular interpolant. c. A non-regular point set having no realization, a non-regular interpolant of this point set and what could be considered an "open" realization of the point set.

a set of samples of an unknown, subjacent object, which is the ultimate goal of shape reconstruction.

By contrast, our regularity condition, which is based on discrete analogues of medial axis and r-sampling, can be checked directly on the sample point set, as the ensuing sections will make clear, and can be turned into a powerful algorithm for reconstructing "desired" connections between adjacent points (we use the word "adjacent" as a synonym of cosimplicial), i.e. for computing a realization of a regular point set. As we shall see, the reconstruction we are able to compute this way for an unorganized and non-regular point set is much better at revealing the inner structure of the point set than the crust. There is a good reason for that: the fact that connections between adjacent points can be recovered depends much more on properties of the point set with respect to these connections than on properties of the subjacent, unknown shape.

3 Properties of regular interpolants

In this section, we prove some fundamental properties of regular interpolants. These properties will be used in Sections 4 and 5 to show that the interpolant of a regular point set can be easily reconstructed from a local analysis of the point set.

The medial axis of an interpolant $\mathcal{O}$ is the Voronoi diagram of $\mathcal{P}$ from which

we have subtracted those components contributed by cosimplicial points of $\mathcal{O}$. Thus, properties of a regular interpolant are intimately linked to properties of its medial axis.

From now on, assume that $\mathcal{O}$ is a regular interpolant of a point set of $\mathbb{R}^m$. We first prove that $\mathcal{O}$ is a subset of the Delaunay graph of $\mathcal{O}$ (also known as the *contour containment condition* [23]).

Proposition 6 (Contour containment) $\mathcal{O} \subset \mathrm{Del}(\mathcal{P})$.

Proof. Let P be a vertex of a simplex S of $\mathcal{O}$ which is assumed not to be Delaunay. Denote by V_i the simplices forming the boundary of the Voronoi region of P (see Figure 8). The contribution of P to the medial axis is defined by the V_i's whose straight-line duals in the Delaunay graph do not belong to $\mathcal{O}$.

For intuition, we first deal with the 2D case. Consider the edge V_i intersecting S. Since S is not in the Delaunay graph, V_i and S are not dual of each other and we thus have the following two possible situations:

(1) (Figure 8.a) V_i is not dual to a Delaunay edge and is thus part of the medial axis. Therefore the thickness $e_{\mathcal{O}}(P)$ at P (distance to the medial axis) is less than the distance from P to the intersection between S and V_i. Since the circumcenter C of S is not inside the Voronoi region of P (otherwise the ball circumscribed to S is void of other sample and S is Delaunay), this distance is necessarily less than or equal to the distance g between P and C which is itself less than or equal to the local granularity $g_{\mathcal{O}}(P)$ at P. Thus:

$$e_{\mathcal{O}}(P) < g \leq g_{\mathcal{O}}(P),$$

violating the regularity hypothesis.

(2) (Figure 8.b) V_i is dual to a Delaunay edge S' of $\mathcal{O}$. Then P does not contribute to the medial axis in the component R' of $\mathbb{R}^2 \setminus \mathcal{O}$ delimited locally by S and S', again violating the regularity condition.

The above proof extends to the 3D case by considering the V_i's intersecting the region H defined by the angular sector of the plane supported by S and delimited by the edges of S incident to P. Again, two situations arise:

(1) The V_i's intersecting H are not all dual to Delaunay simplices. Thus at least one of the V_i's is part of the medial axis and the distance from P to the intersection between H and the medial axis is therefore necessarily less than or equal to the distance to the circumcenter of S, violating the regularity hypothesis.

(2) The V_i's intersecting H are all dual to Delaunay simplices S' of $\mathcal{O}$. Then, P does not contribute to the medial axis inside the region delimited by S and all the Delaunay simplices S', again violating the regularity hypothesis.

Thus, a simplex of a regular interpolant belongs to the Delaunay graph of its vertex set. □

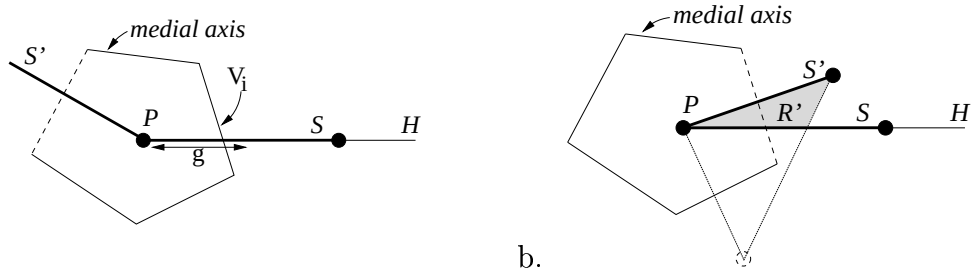

Fig. 8. Illustration of contour containment in the 2D case. The two possible configurations around P when S is not Delaunay: a. the local thickness is smaller than the local granularity; b. P does not contribute to the medial axis in a component of $\mathbb{R}^2 \setminus \mathcal{O}$.

Consequently, the medial axis of $\mathcal{O}$ has the same homological type as $\mathbb{R}^m \setminus \mathcal{O}$ when the shape is regular. Furthermore, simplices of a regular interpolant satisfy the following property which will be directly used in Sections 4 and 5 to characterize and reconstruct interpolants of regular 2D and 3D point sets.

Proposition 7 (Circumcenter on Voronoi region boundary) *Let S be a $(m-1)$-simplex of a regular interpolant $\mathcal{O}$. The circumcenter C of S belongs to the boundary of the Voronoi region of a sample point P of $\mathcal{O}$ if and only if P is a vertex of S.*

Proof. Denote by V_i the simplices forming the boundary of the Voronoi region of a vertex P of S. If the circumcenter C of S belongs to the Voronoi region of a sample point P' which is not a vertex of S, then one of the V_i's is at a distance less than or equal to the granularity at P (see Fig. 9.a) and the situation is analogue to the proof of Prop. 6. Thus, either the medial axis is at a distance less than or equal to the granularity or P does not contribute to the medial axis in a component of $\mathbb{R}^m \setminus \mathcal{O}$ it borders. □

This result is illustrated by Fig. 9.

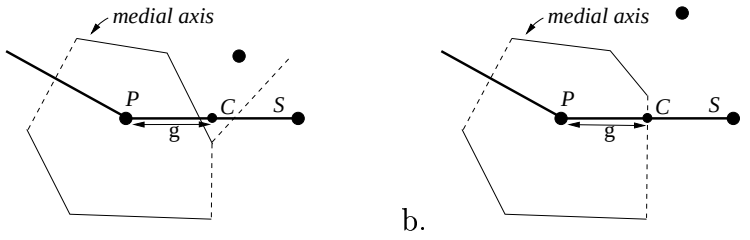

Fig. 9. Illustration of Prop. 7 in the 2D case. The two configurations at P when S is Delaunay: a. $\mathcal{O}$ is not regular at P; b. $\mathcal{O}$ is regular at P.

4 Regular 2D point sets and their interpolants

In this section, we focus on properties of regular point sets of $\mathbb{R}^2$ and their regular interpolants. As we show, the condition of regularity of a point set gives birth to local properties that sample points must satisfy. Given a regular point set, these properties can in turn be used to recover the associated interpolant.

In the rest of this section, $\mathcal{O}$ is assumed to be a regular interpolant of a regular point set $\mathcal{P}$ of $\mathbb{R}^2$.

4.1 Properties

Our aim is to determine what points are adjacent in a regular point set. We will prove here first that two adjacent points on $\mathcal{O}$ define a ball empty of other sample points.

Proposition 8 (Empty ball 2D) *Let P_i and P_j be two adjacent points on $\mathcal{O}$. Then the closed ball B_{ij} with diameter P_iP_j is void of sample points other than P_i and P_j.*

Proof. This is a direct consequence of Prop. 7: the circumcenter of the ball B_{ij} circumscribed to P_iP_j belongs to the boundary of the Voronoi region at P_i and P_j only. Thus, this ball is necessarily void of any other sample point. □

A first consequence of this property is that an edge P_iP_j of a regular interpolant belongs to the *Gabriel graph* (GG) of the point set $\mathcal{P}$, which is included in the Delaunay diagram of $\mathcal{P}$. (In the GG, two vertices P_i and P_j are connected if the ball centered at $(P_i + P_j)/2$ and passing through P_i and P_j contains no other point in its interior.)

Prop. 8 also implies the following result which, while not needed for the ensuing proofs, makes an interesting parallel with Voronoi diagrams. Call *neighborhood*

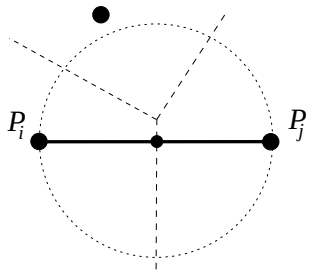

Fig. 10. Adjacent points define an empty ball.

$\mathcal{N}_{\mathcal{O}}(P)$ of a sample point P the set of points of $\mathcal{O}$ having P as their closest sample point. Then:

Proposition 9 (Connectivity of neighborhoods) *The neighborhood of any vertex of a regular interpolant is connected.*

Proof. Consider the situation depicted in Fig. 11. The neighborhood of Q is necessarily between the two perpendicular bisectors of segments PQ and QR. Because of the empty ball property of Prop. 8, the segments pQ and Qr belong to $\mathcal{N}_{\mathcal{O}}(Q)$. Now suppose that $\mathcal{N}_{\mathcal{O}}(Q)$ is not connected, i.e. there is a point x of $\mathcal{O}$ not on pQ and Qr belonging to the neighborhood of Q. The ball centered at x of radius $d(x, Q)$ is thus empty of other vertices of $\mathcal{O}$. But x is a point of $\mathcal{O}$, so is on some edge ST. Clearly, the ball centered at $(S + T)/2$ with diameter $d(S, T)$ contains point Q, thus violating Prop. 8. We have a contradiction. □

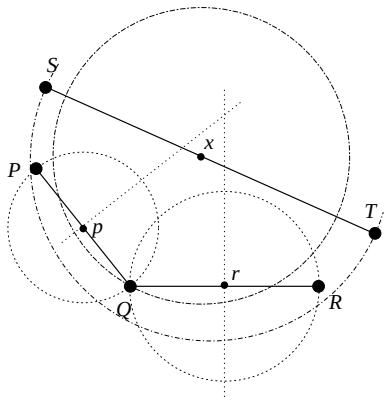

Fig. 11. The neighborhood of a vertex is connected.

Thus, for regular interpolants, the local granularity at a sample P is equivalent to the maximum distance between P and any point in its neighborhood, i.e. it is a measure of the size of the neighborhood.

Proposition 8 tells us that an edge P_iP_j of a regular interpolant $\mathcal{O}$ belongs to the Gabriel graph. But not all edges of the GG correspond to edges of $\mathcal{O}$. The

following property allows to discriminate:

Proposition 10 (Minimal granularities 2D) *The vertices of $\mathcal{P}$ adjacent to a vertex P_i of a regular interpolant $\mathcal{O}$ are the points $P_j, P_k \in \mathcal{P}, P_j \neq P_k$, such that:*

(1) P_iP_j and P_iP_k are edges of the Gabriel graph;
(2) P_j and P_k are the points minimizing the local granularity at P_i (i.e., among the neighbors of P_i in the GG, P_j and P_k are the closest to P_i).

Proof. The situation is as depicted in Fig. 12. Suppose that there is a point P_l such that P_iP_l belongs to the GG of $\mathcal{P}$ and such that P_l is closer to P_i than one of its two adjacent points. Then, the ball centered at M, middle of P_iP_l, is empty (because P_iP_l is in the GG) and thus M is a point of the medial axis of $\mathcal{O}$. Since M is closer to P_i than one of its two adjacent points, the thickness e at P_i is less than the density g at P_i, contradicting the fact that $\mathcal{O}$ is regular. □

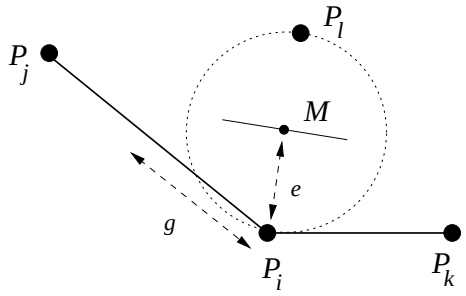

Fig. 12. Points adjacent to P_i are those minimizing the local granularity.

4.2 Reconstruction

Properties of regular interpolants define a theoretical context in which reconstruction of polygonal curves from their sampling points can be easily achieved. The simplest approach for reconstructing a regular interpolant (when it exists) consists in applying Proposition 10 directly. For each point P_i of $\mathcal{P}$, its closest point P_j in $\mathcal{P}$ is determined (the edge P_iP_j necessarily belongs to the GG), giving one edge of the polygonal approximation. Then, the second of the points adjacent to P_i is computed easily, giving Algorithm 1.

The correctness of this algorithm and its ability to identify the desired edges of a regular interpolant follows from Propositions 8 and 10. Note that the algorithm uses the fact that the point P_k which is determined defines necessarily, with the point P_i, an edge of the Gabriel graph of $\mathcal{P}$. This is not true if $\mathcal{P}$ does not come from a regular interpolant (see Section 6).

Algorithm 1 2D reconstruction, regular case

1: **for all** points $P_i \in \mathcal{P}$ **do**
2: **find** the closest point $P_j \in \mathcal{P}$ to P_i
3: **let** $\mathcal{P}'$ be the subset of $\mathcal{P}$ of points P_l such that $P_j \notin B_{il}$
4: **find** the point $P_k = \min_{P_l \in \mathcal{P}'} d(P_i, P_l)$
5: **add** P_iP_k and P_jP_k to the polygonal reconstruction
6: **end for**

Of course, the steps of this incremental construction algorithm do not reflect the way regular interpolant construction is really implemented. Since the regular interpolant is a subset of the Gabriel graph, we can start by computing the GG of the point set $\mathcal{P}$ and then remove from this graph all edges that do not satisfy the second condition of Prop. 10. Since constructing the Delaunay triangulation and the GG of n points can be done in $O(n \log n)$ time, the overall algorithm is also $O(n \log n)$.

An example of reconstruction is shown in Figure 13, along with the Delaunay and Gabriel graphs of the point set.

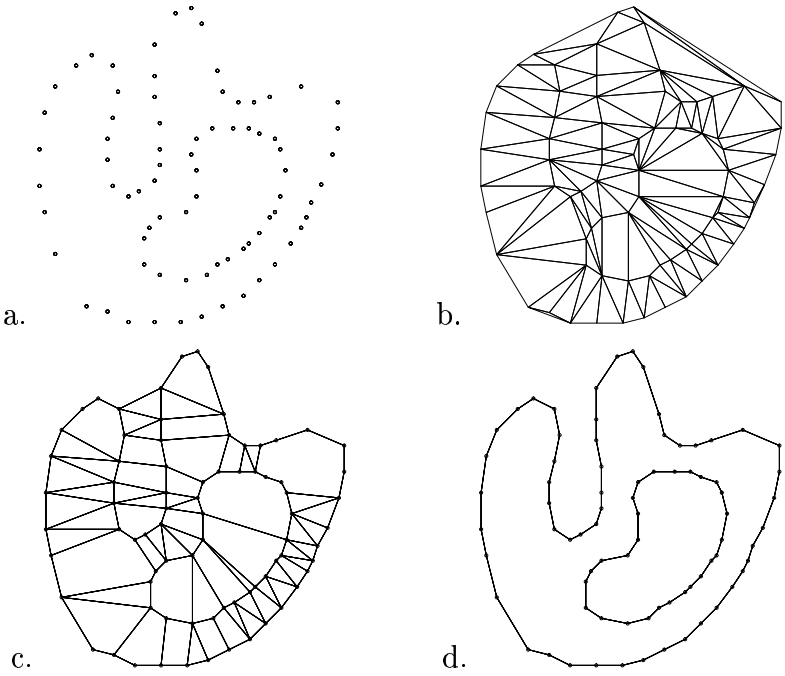

Fig. 13. Example of polygonal reconstruction from a regular sampling of a plane curve. a. Sample points. b. The Delaunay graph. c. The Gabriel graph. d. The regular interpolant.

5 Regular 3D point sets and their interpolants

We now turn to the 3D case and consider properties of adjacent points of regular interpolants of 3D point sets. Since the proofs for the 3D case are very similar to the proofs given in Section 4, we omit them for sake of conciseness. They all follow from Proposition 7.

$\mathcal{O}$ is now assumed to be a regular interpolant of a point set $\mathcal{P}$ of $\mathbb{R}^3$.

5.1 Properties

Proposition 11 (Empty ball 3D) *Let P_i, P_j and P_k be adjacent points on $\mathcal{O}$. Then the closed ball B_{ijk} with center and radius those of the circle circumscribed to P_i, P_j and P_k is void of sample points other than P_i, P_j and P_k.*

As in the 2D case, Proposition 11 implies that the neighborhood of a vertex of a regular interpolant is connected. More importantly, it implies that $\mathcal{O}$ is a subcomplex of a suitable extension of the Gabriel graph to 3D.

Let us define this extension of the Gabriel graph of a point set as follows. Let P_i, P_j, P_k be three points of $\mathcal{P}$. Consider the simplicial complex formed as follows: if the closed ball with center and radius those of the circle circumscribed to P_i, P_j, P_k contains no other sample points than P_i, P_j and P_k, then add edges between P_iP_j, P_jP_k, P_kP_i and add a 2-simplex (a triangle) with vertices $P_iP_jP_k$. For lack of a better name (this object does not seem to have been studied in the literature), we call this complex the *3D Gabriel complex* (3DGC) of $\mathcal{P}$. It is intuitively clear that the 3DGC is a subcomplex of the Delaunay diagram of $\mathcal{P}$.

A consequence of Proposition 11 is that $\mathcal{O}$ is a subcomplex of 3DGC. But not all 2-simplices of the 3DGC of $\mathcal{P}$ correspond to triangles of $\mathcal{O}$. An argument similar to the 2D case shows that:

Proposition 12 (Minimal granularities 3D) *Let P_i, P_j be two adjacent points of a regular interpolant $\mathcal{O}$ of $\mathcal{P}$. The sample points P_k and P_l forming adjacent triples with P_i and P_j are those points of $\mathcal{P}$ such that:*

(1) the triangles $P_iP_jP_k$ and $P_iP_jP_l$ belong to the 3DGC;
(2) the local granularities at P_i and P_j are minimal, i.e., points P_k and P_l are those vertices of the 3DGC which minimize the circumradius at P_i and P_j.

Based on the above properties, the algorithm for reconstructing the interpolant of a regular sample point set is very similar to the planar case. The idea is to follow directly Proposition 12.

Let $E(P_i, P_j, P_k)$ denote the edge (P_i, P_j) associated to the triangle (P_i, P_j, P_k). Algorithm 2 is used to reconstruct the interpolant.

Algorithm 2 3D reconstruction, regular case

find an initial adjacent triple P_0, P_1, P_2
initialize list *Ledges* with $E(P_0, P_1, P_2)$, $E(P_1, P_2, P_0)$ and $E(P_0, P_2, P_1)$
for all edges $E(P_i, P_j, P_k)$ in *Ledges* **do**
  let $\mathcal{P}'$ be the subset of $\mathcal{P}$ of points P_m such that $P_k \notin B_{ijm}$
  find the point $P_l = \min_{P_m \in \mathcal{P}'} \text{radius}(B_{ijm})$
  add the triangle (P_i, P_j, P_l) to the triangular mesh
  if $E(P_i, P_l, P_j)$ and $E(P_j, P_l, P_i)$ are not in *Ledges* **then**
    add them to the list
  else
    remove them
  end if
  remove $E(P_i, P_j, P_k)$ from *Ledges*
end for

An example of the use of this algorithm for reconstructing the interpolant of points randomly distributed on a sphere is shown in Figure 14.

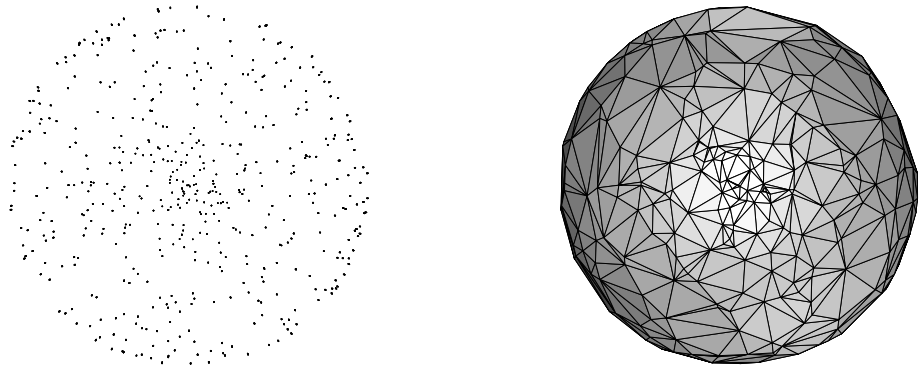

Fig. 14. Example of polygonal reconstruction from points on a sphere. In this example, points are randomly distributed on a sphere. The local thickness is therefore constant over the set of points and equal to the sphere radius.

Note that

- To find an initial adjacent triple, we can proceed as follows: choose a point P_0 in $\mathcal{P}$, find its nearest neighbor P_1 and pick the point P_2 minimizing

the radius of the ball B_{012}. This ball is necessarily empty of other samples. Indeed, if a sample P_k is inside B_{012} and such that the radius of B_{01k} is larger than the radius of B_{012}, then P_k is closer to P_0 than P_1.

- If $\mathcal{O}$ is not connected, i.e., is composed of more than one component, then Algorithm 2 must be iterated.
- An important difference with the 2D case is that here reconstruction can only be carried out by an incremental surface construction algorithm, and not by elimination from the Delaunay complex. In 3D, three sample points can satisfy Proposition 12 without the 2-simplex linking the three points being part of a regular interpolant. An advancing-front technique is thus the only alternative in this case.

6 Reconstructing non-regular point sets

In the previous sections, we have seen how to reconstruct the interpolant of a point set when the points are regularly distributed. Now, regularity is a somewhat restrictive notion: even when it is dense, a mesh may not be everywhere regular depending on local curvature properties of the unknown subjacent surface. Consider first a set of points uniformly and randomly distributed on a sphere. The interpolant of these samples is regular since the thickness is almost a constant, equal to the sphere radius. By contrast, for points randomly distributed on the surface of a torus, then whatever the density there may be regions of the interpolant which are not regular. This is due to the fact that the tangent plane at hyperbolic samples intersects the mesh locally, so there is a tendency to have flat four-point configurations in the neighborhood of such samples.

The difficulty of finding regular point sets of negative-curvature surfaces leads us to introduce in this section the notion of minimal interpolant. Using techniques which closely parallel the regular case, we propose an approach for reconstructing non-regular point sets that leads to polygonal structures which describe well the adjacency between points. Such structures can then be used to extract interpolating curves (in 2D) or surfaces (in 3D) which are manifold.

6.1 Minimal interpolants

With Propositions 10 (in 2D) and 12 (in 3D), we have seen simple criteria for iteratively constructing the interpolant of a regular point set. The idea was to determine, from an already computed $(m-1)$-simplex s_1, the neighboring $(m-1)$-simplex s_2 such that:

- s_1 and s_2 are locally valid (i.e., the ball circumscribed to s_1 does not contain s_2 and conversely) ;
- the granularities at the vertices of the $(m-2)$-simplex connecting s_1 and s_2 are minimal.

When the point set $\mathcal{P}$ is regular, s_1 and s_2 necessarily belong to the Gabriel complex (graph in 2D) of $\mathcal{P}$.

We want to apply a similar procedure to non-regular point sets. There is, however, an important difference: the $(m-1)$-simplices thus computed do not necessarily belong to the Gabriel complex of $\mathcal{P}$. This behavior is not desirable since we would like interpolants of general point sets to share the properties of interpolants of regular point sets. For this reason, we use the following heuristics for reconstructing general point sets:

(1) the $(m-1)$-simplex s_2 neighboring an already computed $(m-1)$-simplex s_1 on $\mathcal{O}$ is such that the granularities at the vertices of $s_1 \cap s_2$ are minimal;
(2) s_1 and s_2 belong to the Gabriel complex of $\mathcal{P}$.

In the case of regular point sets, these two conditions are equivalent to Propositions 10 and 12. They are however more restrictive in the general case and eliminate some of the points that are potentially adjacent to a point in 2D or an edge in 3D. The resulting approximation corresponds to a subcomplex of the Gabriel complex of $\mathcal{P}$ which we call a *minimal interpolant* of $\mathcal{P}$. Note that, in general, minimal interpolants are not closed as in the regular case (see Fig. 17.b).

The algorithms we derive from the above heuristics are very similar to the regular case. For instance, in 3D, and keeping the same notations as before, we use Algorithm 3.

Algorithm 3 3D reconstruction, non-regular case

1: **find** an initial adjacent triple P_0, P_1, P_2
2: **initialize** list *Ledges* with $E(P_0, P_1, P_2)$, $E(P_1, P_2, P_0)$ and $E(P_0, P_2, P_1)$
3: **for all** edges $E(P_i, P_j, P_k)$ in *Ledges* **do**
4: **let** $\mathcal{P}'$ be the subset of $\mathcal{P}$ of points P_m / $P_k \notin B_{ijm}$
5: **find** the point $P_l = \min_{P_m \in \mathcal{P}'} \text{radius}(B_{ijm})$
6: **if** the triangle (P_i, P_j, P_l) is in the 3DGC and not already in $\mathcal{O}$ **then**
7: **add** (P_i, P_j, P_l) to the triangular mesh $\mathcal{O}$
8: **add** $E(P_i, P_l, P_j)$ and $E(P_j, P_l, P_i)$ to *Ledges*
9: **end if**
10: **remove** $E(P_i, P_j, P_k)$ from *Ledges*
11: **end for**

Underlying our reconstruction algorithm in the non-regular case is the implicit idea that the Gabriel complex contains enough triangles for reconstruction.

Why is it so? In the limit of a dense sampling of a smooth surface F, the surface is relatively flat on the scale of the 3DGC triangles that lie along F. So the 3DGC around a vertex P is sure to contain all the triangles that belong to the 2D Delaunay triangulation of the projections of the sample points in a neighborhood of P onto a place locally approximating F.

6.2 Manifold interpolants

In the case of a non-regular point set, the minimal interpolant is not generally manifold. It is rather composed of several manifold components which are connected one to the other. However, in applications where post-processing of the resulting shape is considered (mesh simplification for example), a manifold is usually required. One way to proceed is to first identify all the components and then either to consider them as separate shapes which are connected or to eliminate some of them until there is only one manifold component.

To identify the different components of a minimal interpolant, groups of simply connected simplices are extracted. This is done by searching edges having each of their endpoints shared by a unique edge in 2D and triangles having each of their edges shared by a unique triangle in 3D, and grouping them (see Figure 15).

The next step is either to keep separate components or to proceed by elimination. Even though separate components are satisfactory in some applications, they often lead to fragmented interpolants which can hardly be handled. The alternative is thus to consider an elimination step. To this purpose, several methods can be applied depending on the heuristic that is chosen. We use an elimination step which consists in iteratively eliminating groups with small numbers of simplices and recomputing groups until there is only one component (see Figure 15). Groups which have the same number of simplices are compared by their lengths in 2D and their areas in 3D, the smallest values being first eliminated.

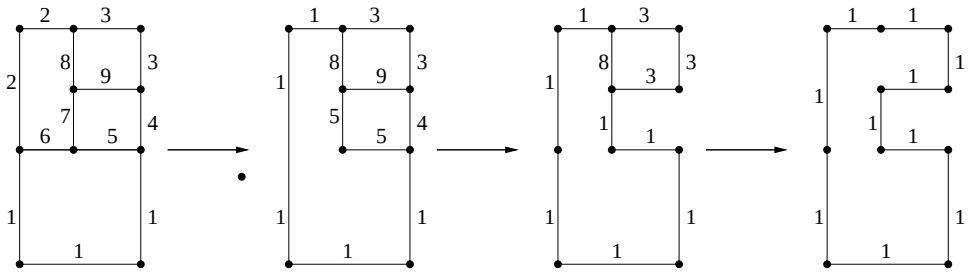

Fig. 15. Grouping and iteratively eliminating groups with small number of vertices.

Such an approach gives satisfactory results in most situations as shown in the results (§ 7.2). However, and since this method relies on one heuristic which

is that small groups of vertices with small sizes (length or area) should not appear in the final shape, we are also currently investigating other approaches based on different heuristics.

7 Implementation and results

The reconstruction methods presented in this paper have been implemented and throughly tested on different types of data sets. In this section, we elaborate on the complexity of the algorithms used and report the results of our tests.

7.1 Complexity

The algorithms given in the previous sections are iterative and consist in determining neighboring simplices of already computed ones. Therefore, complexities of such algorithms mostly depend on the search functions used.

In the regular case, search functions aim at finding the vertex which minimizes the granularity. In their simplest form (Algorithms 1 and 2), these functions search for a vertex over the whole set of sample points and therefore have an asymptotic complexity of $O(n)$, where n is the number of sample points. This complexity can easily be reduced by using appropriate data structures. Examples of such data structures are k-d trees and Voronoi diagrams [27,28]. Let $O(f(n))$ be the asymptotic complexity of the search function which depends on the space dimension and the data structure which is used. Then, since the number of simplices of a regular interpolant is linear in n, the algorithms have asymptotic complexities of the form: $O(nf(n))$. For example in the 2D case, the algorithm takes $O(n \log n)$ by using the Voronoi diagram as a search structure and such an algorithm is furthermore optimal in terms of complexity [28].

The situation is more complex for non-regular point sets (Algorithm 3). In that case, the search function still aims at finding the vertex which minimizes the granularity, but the fact that the resulting simplex belongs to the Gabriel graph must be verified. The latter task, which consists in verifying the emptiness of the ball circumscribed to the simplex, is a classical geometric range searching problem. Again, the query time depends on the data structure which is used (see [29,30] for reviews on geometric range searching problems). In the 2D case and again by using the Voronoi diagram, a search will take time $O(\log n)$. In our implementation, we use a linear-size data structure which leads to $O(n^{1-1/m})$ search functions for points uniformly distributed in $\mathbb{R}^m$.

Let $O(f(n))$ be the asymptotic complexity of the search function. Then, the asymptotic complexity of the algorithm is $O(Nf(n))$, where N is the number of simplices in the interpolant. N is generally linear in n but can be larger than that depending on the natural organization of the point set. It should be noted here that using the Voronoi diagram (or equivalently the Delaunay complex) as a pre-processed data structure does not necessarily improve the complexity and the running time in the 3D case because of the pre-processing time required.

Figure 16 shows running times of Algorithm 3 on a Sparc Ultra 30 workstation (270 MHz). Sets of points are randomly distributed on a torus (see Figure 18 for reconstruction examples). It takes 0.61 second to handle 500 points and approximately 33 minutes to handle 100,000 points. The asymptotic complexity is "approximately" $O(n^{3/2})$ as shown by the dotted curve. This can be explained by the fact that points are, in that case, uniformly distributed on a surface and not in $\mathbb{R}^3$. Therefore, the search function is "approximately" $O(n^{1/2})$ (and not $O(n^{2/3})$) and consequently the algorithm runs in $O(n^{3/2})$ since the number of simplices on the torus is $2n$ (by the Euler formula).

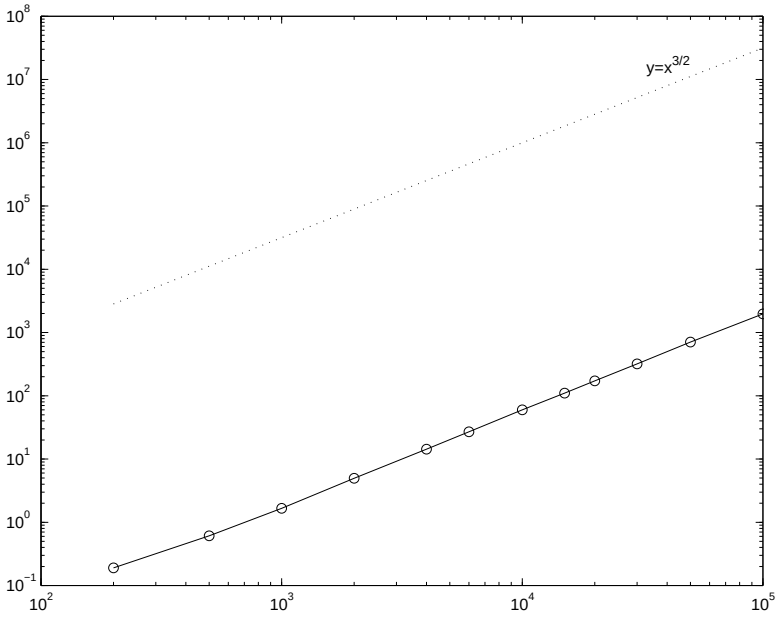

Fig. 16. Log-log scale plots of the running time function for points randomly distributed on a torus and for the function $y = x^{3/2}$ (dotted curve).

7.2 Results

We now show some results of the application of the reconstruction algorithms presented in this paper to non-regular point sets. In 2D, Fig. 17.a and .b show the minimal interpolant of a set of points sampled from a closed curve. This non-regular point set should be compared to the regular point set of Fig. 13.a.

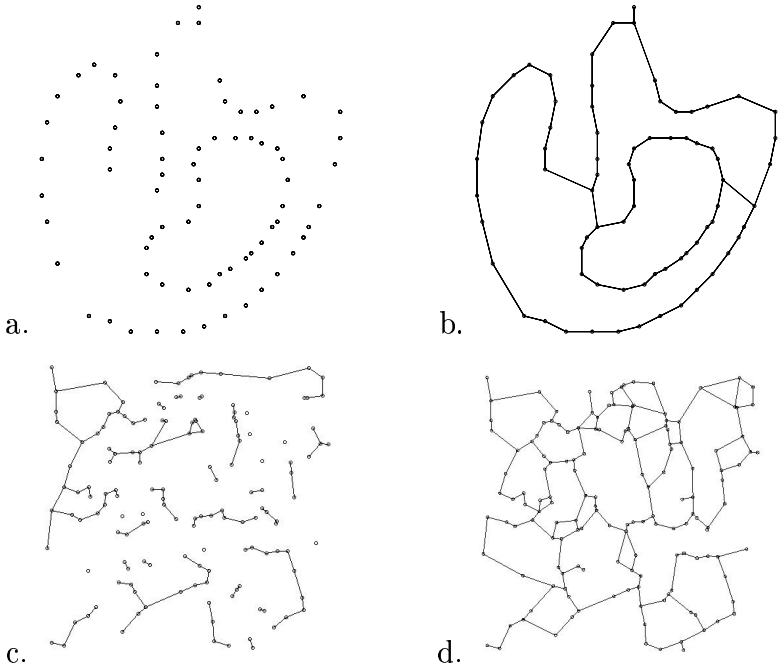

Fig. 17. Interpolants of non-regular planar point sets. a. Sample points on a curve (compare with Fig. 13). b. The minimal interpolant computed. c. The crust [21] of a set $\mathcal{P}$ of randomly distributed points in the plane. d. The minimal interpolant of $\mathcal{P}$. Note that the minimal interpolant reveals more adjacency relations than the crust.

Because of the poor sampling of the curve at some places, it is not possible to reconstruct a shape homeomorphic to the original curve by considering local connections. Fig. 17.c and .d compare the interpolant reconstructed with our algorithm to the shape obtained with the crust algorithm [21]. Note how the adjacency relations computed by our algorithm reveal the natural organization of the point set. Note also how the minimal interpolant shows more of the internal structure of the point set than the crust.

In 3D, Fig. 18 shows the triangular mesh obtained when reconstructing the minimal interpolant of a set of randomly distributed points on a torus. Note that by contrast to the case of points randomly distributed on the sphere, points on a torus do not, in general, define a regular interpolant.

The remaining examples show the minimal interpolants computed from data of various origin. These examples illustrate the ability of the algorithm to adapt to different types of data. Fig. 19 considers reconstruction of point sets obtained by reconstruction from image sequences using occluding contours [31]. Fig. 20 deals with a point set representing a dragon that was obtained from the 3DSite on the Internet (`http://www.3dsite.com`). Fig. 21 shows the reconstruction of a mechanical part taken from Hoppe's web site [10]. Finally, Fig. 22 shows the reconstruction of a terrain model taken from the site of the Swiss Federal Office of Topography (`http://www.swisstopo.ch`).

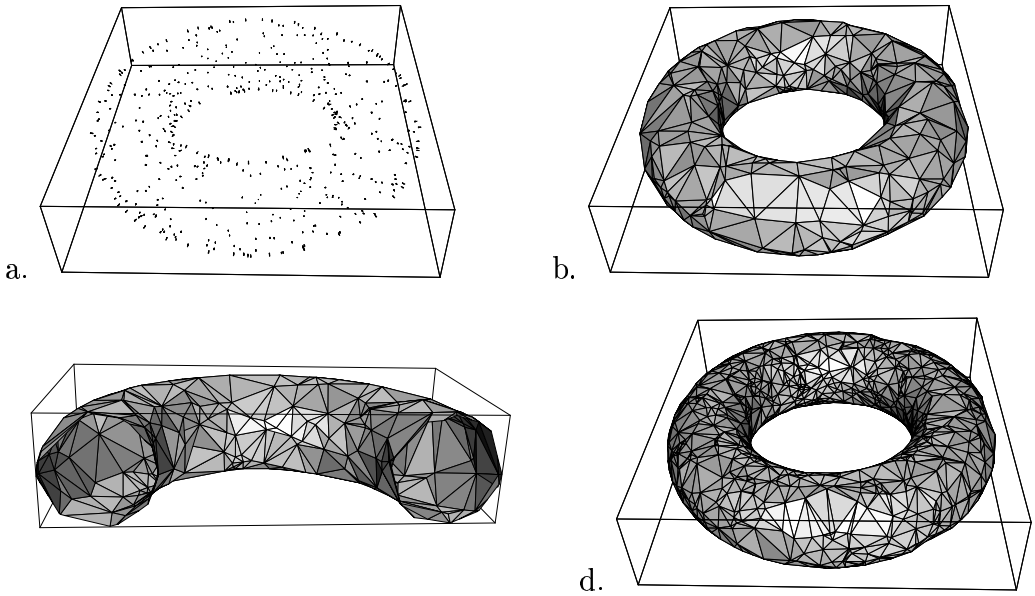

Fig. 18. Reconstructing points on a torus. a. 500 points randomly distributed on a torus. b. The manifold interpolant of a. c. A cross section of the reconstruction. d. Reconstruction of 1,000 points randomly distributed on a torus.

8 Conclusion

In this paper, we have introduced the notion of regular interpolant of a 2D or 3D point set. From the properties of regular interpolants, reconstruction methods have been proposed for regular and non-regular sets of points which are intermediate between spatial subdivision algorithms and advancing-front techniques. When the point set considered is not regular, heuristics have been proposed to extract a manifold interpolant. The resulting reconstruction methods lead to very simple and efficient algorithms that iteratively build interpolants of sets of points. These algorithms have all been implemented and validated on a wide range of input data.

We are currently working on further improvements and applications of our method. In particular, we are working on the following aspects:

- In the current version of our algorithm, the final step is simple and consists in extracting connected components. This part of the process can be improved by considering other heuristics.
- We are investigating more closely what it means for a point set to be regular and how to ensure that a point set will have a regular interpolant. This is a crucial issue for applications. We are also investigating further theoretical implications of regular interpolants. It would be very interesting to know precisely how easy, or hard, it is for a point set to be regular.
- Another avenue we consider is the problem of designing sampling strategies which would ensure that the resulting point set is regular. Such strategies

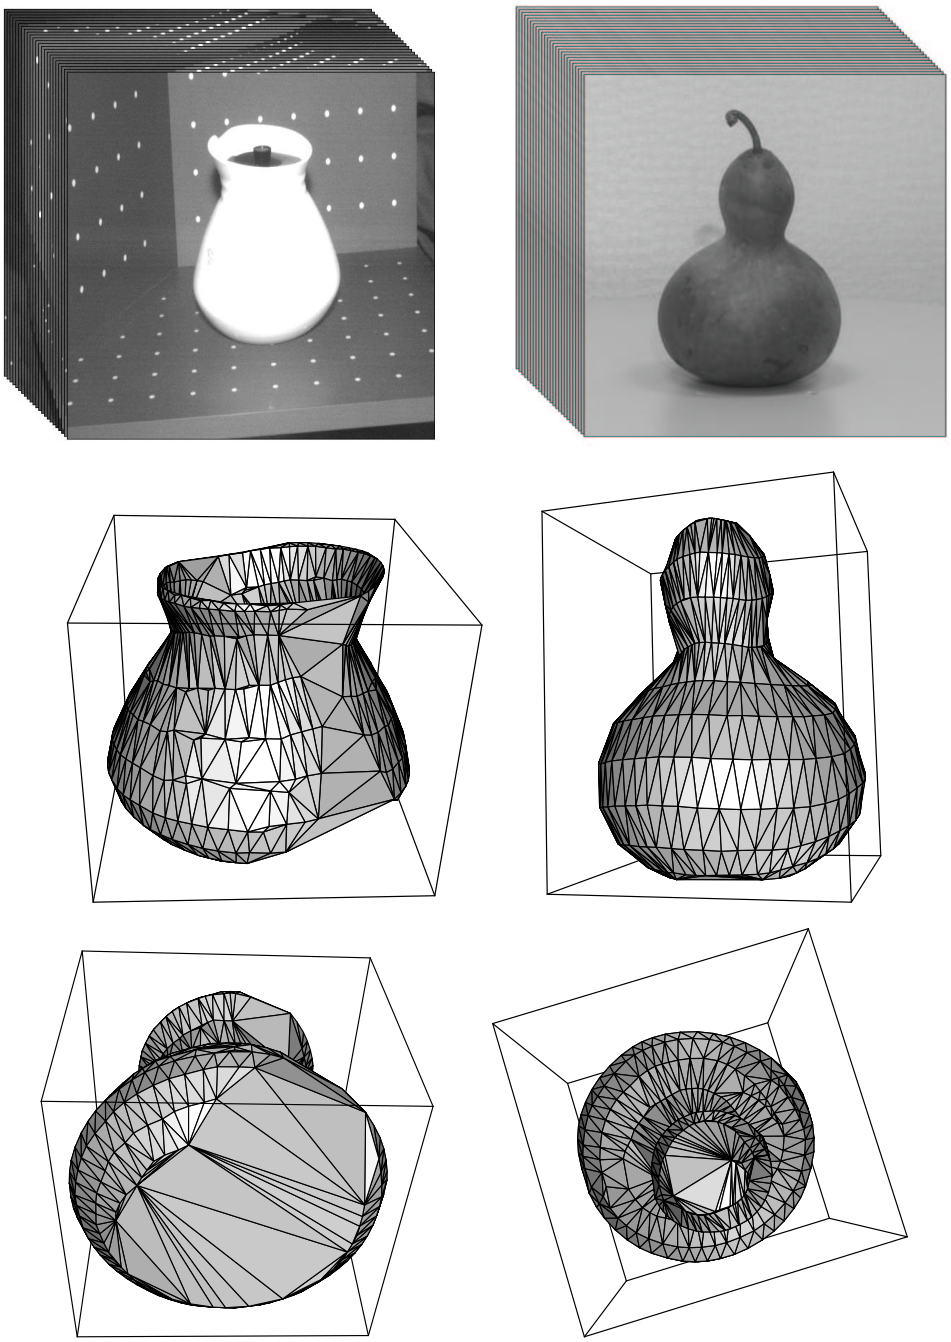

Fig. 19. Manifold interpolants of point sets obtained by reconstruction from image sequences [31]. Running times on a Sparc Ultra 30 workstation are 0.9 s for the teapot (540 points) and 1.05 s for the calabash (571 points).

may provide guidance to individuals planning range scans. For this, we need to look more closely at the relationships between discrete and continuous medial axes and how our sampling criterion relates to the lfs criterion.

- We want to study further properties and applications of medial axes of interpolants as we defined them. A fairly interesting parallel between discrete and continuous medial axes can be made by considering the discrete medial

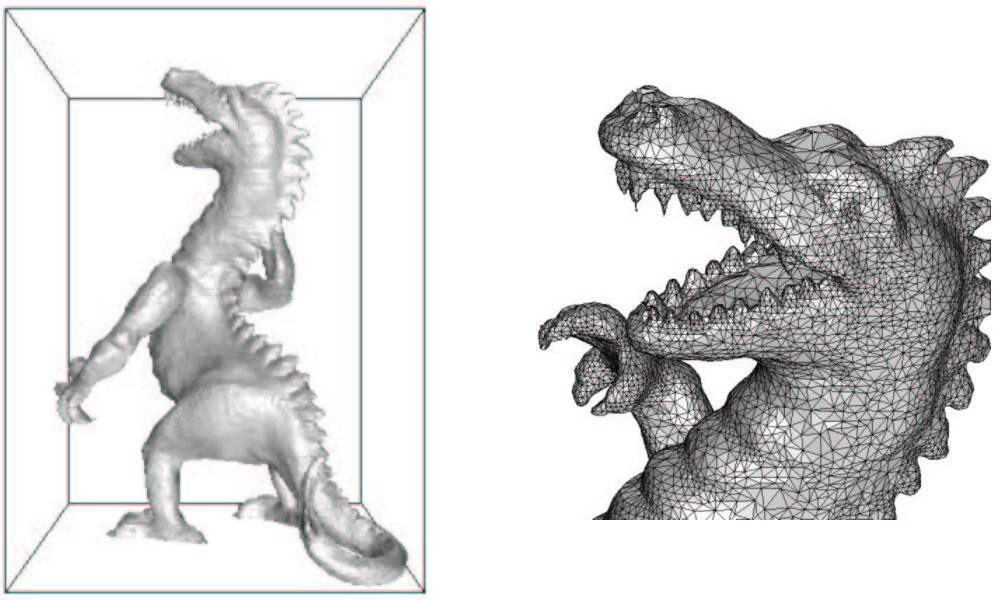

Fig. 20. Reconstruction of the dragon model (25,000 points). Running time is 3 minutes 55 seconds on a Sparc station.

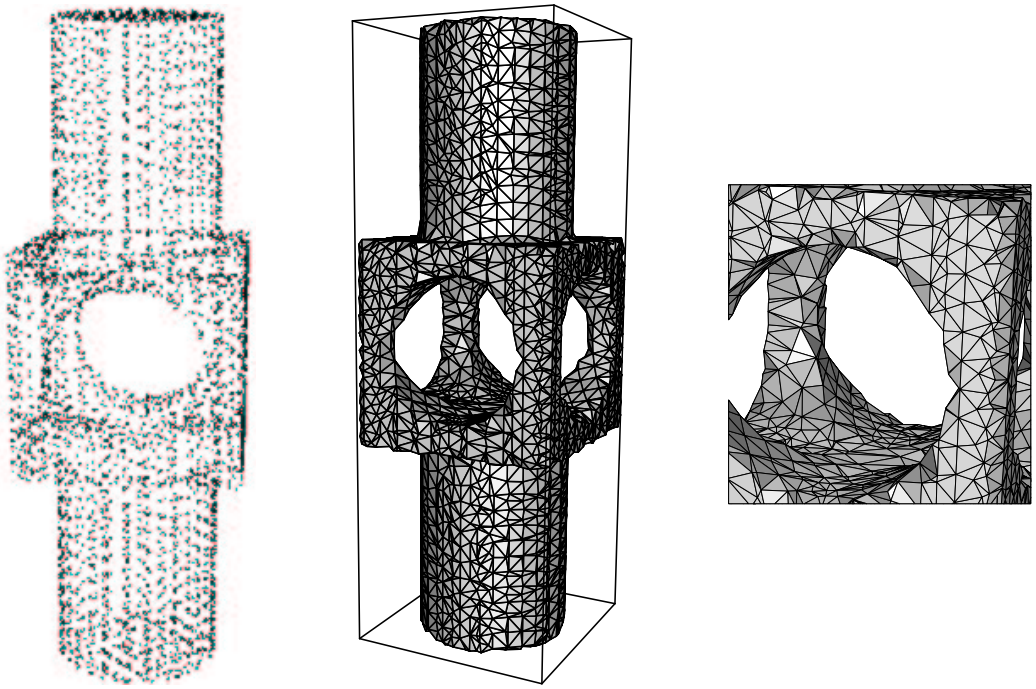

Fig. 21. Reconstruction of a CAD model (taken from Hoppe's web site [10]). Running time is 16 seconds (4,105 points).

axis as the continuous medial axis of a union of maximal balls built over the interpolant. Medial axes of unions of balls have been studied by Attali [26] and more recently by Amenta and Kolluri [32], who intend to use them as tools for polygonal surface reconstruction. Clearly, our research has a lot in common with the work of these authors and we intend to explore the

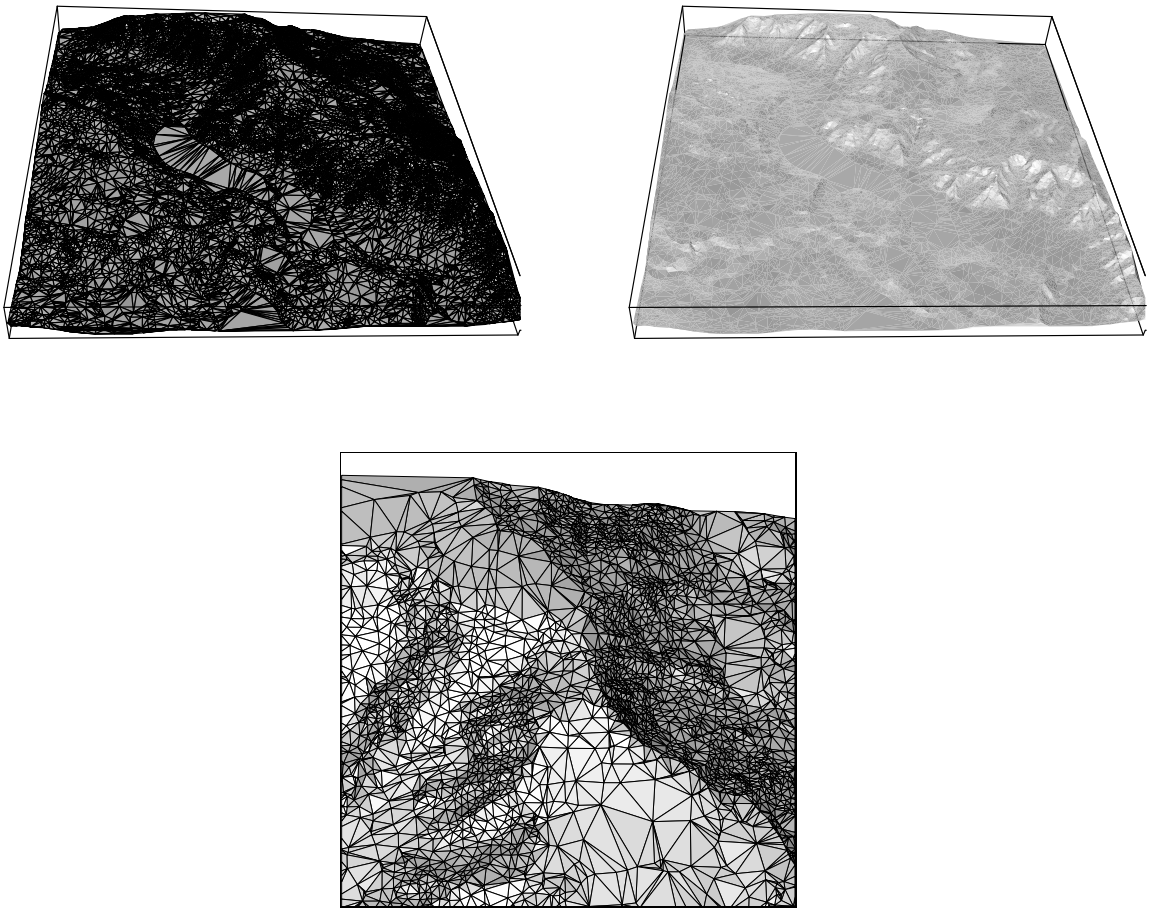

Fig. 22. Reconstruction of a terrain model (17,000 points). This test model was taken from the site of the Swiss Federal Office of Topography (`http://www.swisstopo.ch`). Points are sampled on the vectorization of the contour lines of a terrain map. Running time is 3 minutes and 18 seconds for this point set.

connections in the near future.

References

[1] E. Mencl, H. Müller, Interpolation and approximation of surfaces from three-dimensional scattered data points, in: State of the Art Reports, Eurographics'98, 1998, pp. 51–67.

[2] J.-D. Boissonnat, Geometric structures for three-dimensional shape reconstruction, ACM Transactions on Graphics 3 (4) (1984) 266–286.

[3] R. Veltkamp, Boundaries through scattered points of unknown density, Graphical Models and Image Processing 57 (6) (1995) 441–452.

[4] H. Edelsbrunner, E.-P. Mücke, Three-dimensional alpha shapes, ACM Transactions on Graphics 13 (1) (1994) 43–72.

[5] B. Guo, J. Menon, B. Willette, Surface reconstruction using alpha shapes, Computer Graphics Forum 16 (4) (1997) 177–190.

[6] D. Attali, r-regular shape reconstruction from unorganised points, Computational Geometry: Theory and Applications 10 (4) (1998) 239–247.

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

[8] N. Amenta, M. Bern, Surface reconstruction by Voronoi filtering, Discrete and Computational Geometry 22 (4) (1999) 481–504.

[9] N. Amenta, S. Choi, T. Dey, N. Leekha, A simple algorithm for homeomorphic surface reconstruction, in: Proc. of SCG'00 (ACM Symposium on Computational Geometry), 2000, pp. 213–222.

[10] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, W. Stuetzle, Surface reconstruction from unorganized points, in: ACM Computer Graphics (Proceedings SIGGRAPH), Vol. 26, 1992, pp. 71–78.

[11] B. Curless, M. Levoy, A volumetric method for building complex models from range images, in: ACM Computer Graphics (Proceedings SIGGRAPH), 1996, pp. 303–312.

[12] F. Bernardini, C. Bajaj, J. Chen, D. Schikore, Triangulation-based object reconstruction methods, in: Proc. of SCG'97 (ACM Symposium on Computational Geometry), 1997, pp. 481–484.

[13] J.-D. Boissonnat, F. Cazals, Smooth surface reconstruction via natural neighbor interpolation of distance functions, in: Proc. of SCG'00 (ACM Symposium on Computational Geometry), 2000, pp. 223–232.

[14] H.-K. Zhao, S. Osher, B. Merriman, M. Kang, Implicit, nonparametric shape reconstruction from unorganized points using a variational level set method, to appear in Computer Vision and Image Understanding (2000).

[15] E. Mencl, H. Müller, Graph-based surface reconstruction using structures in scattered point sets, in: Proc. of CGI'98 (Computer Graphics International), 1998, pp. 298–311.

[16] F. Bernardini, J. Mittelman, H. Rushmeier, C. Silva, G. Taubin, The ball-pivoting algorithm for surface reconstruction, IEEE Transactions on Visualization and Computer Graphics 5 (4) (1999) 349–359.

[17] P. Crossno, E. Angel, Spiraling edge: fast surface reconstruction from partially organized sample points, in: IEEE Visualization'99, 1999, pp. 317–324.

[18] M. Gopi, S. Krishnan, A fast and efficient projection-based approach for surface reconstruction, to appear in International Journal of High-Performance Computer Graphics, Multimedia and Vizualisation (2000).

[19] O. Faugeras, Three-Dimensional Computer Vision: A Geometric Viewpoint, Artificial Intelligence, MIT Press, Cambridge, 1993.

[20] F. Bernardini, C. Bajaj, Sampling and reconstructing manifolds using α-shapes, in: Proceedings of 9th Canadian Conference on Computational Geometry, 1997, pp. 193–198.

[21] N. Amenta, M. Bern, M. Kamvysselis, A new Voronoi-based surface reconstruction algorithm, in: ACM Computer Graphics (Proceedings SIGGRAPH), 1998, pp. 415–421.

[22] J. Serra, Image Analysis and Mathematical Morphology, Volume I, Academic Press, 1982.

[23] J.-D. Boissonnat, B. Geiger, Three Dimensional Reconstruction of Complex Shapes Based on the Delaunay Triangulation, Rapport de Recherche 1697, INRIA (Apr. 1992).

[24] R. Ogniewicz, M. Ilg, Voronoi skeletons: Theory and applications, in: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, Urbana Champaign, (USA), 1992, pp. 63–69.

[25] J. Brandt, Convergence and continuity criteria for discrete approximations of the continuous planar skeleton, Computer Vision, Graphics and Image Processing 59 (1) (1994) 116–124.

[26] D. Attali, Squelettes et graphes de Voronoi 2D et 3D, Ph.D. thesis, Université Joseph Fourier - Grenoble I (1995).

[27] J. Bentley, Multidimensional binary search trees used for associative searching, Communications of the ACM 18 (1975) 509–517.

[28] F. Preparata, M. Shamos, Computational Geometry, An Introduction, Springer-Verlag, Berlin, 1985.

[29] J. Matoušek, Geometric Range Searching, ACM Computing Surveys 26 (4) (1994) 421–461.

[30] P.-K. Agarwal, J. Erickson, Geometric Range Searching and its Relatives, Contemporary Mathematics 223, American Mathematical Society Press (1999) 1–56.

[31] E. Boyer, M.-O. Berger, 3D surface reconstruction using occluding contours, International Journal of Computer Vision 22 (3) (1997) 219–233.

[32] N. Amenta, R. Kolluri, Accurate and efficient unions of balls, in: Proc. of SCG'00 (ACM Symposium on Computational Geometry), 2000, pp. 119–128.