



Reconstruction d'ensembles compacts 3D

Frédéric Cazals, David Cohen-Steiner

► To cite this version:

Frédéric Cazals, David Cohen-Steiner. Reconstruction d'ensembles compacts 3D. [Research Report] RR-6868, INRIA. 2009, pp.20. inria-00370208

HAL Id: inria-00370208

<https://inria.hal.science/inria-00370208>

Submitted on 23 Mar 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Reconstructing 3D compact sets

Frédéric Cazals — David Cohen-Steiner

N° 6868

March 2009

Thèmes SYM et BIO

 ***Rapport
de recherche***

Reconstructing 3D compact sets

Frédéric Cazals ^{*}, David Cohen-Steiner [†]

Thèmes SYM et BIO — Systèmes symboliques et Systèmes biologiques
Projet ABS & Geometrica

Rapport de recherche n° 6868 — March 2009 — 17 pages

Abstract: Reconstructing a 3D shape from sample points is a central problem faced in medical applications, reverse engineering, natural sciences, cultural heritage projects, etc. While these applications motivated intense research on 3D surface reconstruction, the problem of reconstructing more general shapes hardly received any attention. This paper develops a reconstruction algorithm changing the 3D reconstruction paradigm as follows.

First, the algorithm handles general shapes i.e. compact sets as opposed to surfaces. Under mild assumptions on the sampling of the compact set, the reconstruction is proved to be correct in terms of homotopy type. Second, the algorithm does not output a single reconstruction but a nested sequence of *plausible* reconstructions. Third, the algorithm accommodates topological persistence so as to select the most stable features only. Finally, in case of reconstruction failure, it allows the identification of under-sampled areas, so as to possibly fix the sampling.

These key features are illustrated by experimental results on challenging datasets, and should prove instrumental in enhancing the processing of such datasets in the aforementioned applications.

Key-words: 3D reconstruction, surface reconstruction, distance function, Voronoi diagram, Morse-Smale complex, flow complex, topological persistence.

^{*} INRIA Sophia-Antipolis-Méditerranée, Algorithms-Biology-Structure; Frederic.Cazals@sophia.inria.fr

[†] INRIA Sophia-Antipolis-Méditerranée, Geometrica; David.Cohen-Steiner@sophia.inria.fr

Reconstruction d'ensembles compacts 3D

Résumé : Reconstruire un modèle à partir d'échantillons est un problème central se posant en médecine numérique, en ingénierie inverse, en sciences naturelles, etc. Ces applications ont motivé une recherche substantielle pour la reconstruction de surfaces, la question de la reconstruction de modèles plus généraux n'ayant pas été examinée. Ce travail présente un algorithme visant à changer le paradigme de reconstruction en 3D comme suit.

Premièrement, l'algorithme reconstruit des formes générales—des ensembles compacts et non plus des surfaces. Sous des hypothèses appropriées, nous montrons que la reconstruction a le type d'homotopie de l'objet de départ. Deuxièmement, l'algorithme ne génère pas une seule reconstruction, mais un ensemble de reconstructions plausibles. Troisièmement, l'algorithme peut être couplé à la persistance topologique, afin de sélectionner les traits les plus stables du modèle reconstruit. Enfin, en cas d'échec de la reconstruction, la méthode permet une identification aisée des régions sous-échantillonnées, afin éventuellement de les enrichir.

Ces points clefs sont illustrés sur des modèles difficiles, et devraient permettre de mieux tirer parti de leurs caractéristiques dans les applications sus-citées.

Mots-clés : Reconstruction de formes en 3D, reconstruction de surfaces, fonction distance, digramme de Voronoi, diagramme de Morse-Smale, *flow complex*, persistance topologique.

1 Introduction

Reconstructing models from samples. Reconstruction is the generalisation of the *connect-the-dots* problem: given a sampling of an unknown model, provide a plausible reconstruction of this model from the samples. Since a variety of devices capturing points on/within a 3D model exist, such as laser range scanners, X-ray, CT or MRI machines, reconstruction has countless applications, namely that of organs in medicine, of computer models (spare parts, mock-ups) in reverse engineering, of art pieces in cultural heritage projects, of plants in natural sciences, etc.

The reconstruction should match the model in terms of geometric and topological properties. To ease the process, assumptions on the geometry and/or topology of the model may be used. Such a-priori do alleviate the reconstruction problem, which yet remains challenging for two reasons. First, the samples may not comply with the hypothesis made on the model. In particular, they might be too sparse for the reconstruction to successfully capture challenging details such as thin parts, holes or boundaries. Second, even for a sampling compliant with these features, a satisfactory reconstruction may not be unique.

For shapes in the ambient 3D Euclidean space, a standard assumption is that the model is a smooth surface. Algorithms reconstructing surfaces report a smooth model or a combinatorial representation, typically a triangle mesh. These two categories have opposite pros and cons, since the latter corresponds to approaches offering flexibility to represent local features, such as holes and boundaries, while the former features strategies returning more compact yet globally defined surfaces. Example algorithms in the first realm are those based on level sets [ZOF01], radial basis functions [CBC⁺01] and moving least squares [ABCO⁺01]. As surveyed in [CG06], examples in the second area are intimately related to the Voronoi diagram of the samples, and its dual the Delaunay triangulation [AB99, GJ02, BC02, ACDL00, Cha03, CSD04, Ede04].

Three-dimensional Reconstruction based upon the flow complex. Because reconstruction boils down to establishing neighborhood connections between samples, geometric complexes encoding proximity relationships may be used as a background. While it has long been recognized that the Delaunay triangulation contains, under mild assumptions, a satisfactory reconstruction [Boi84], it has recently been shown that the *flow complex* also does so. Prosaically, the flow complex consists of the watersheds associated to the critical points of the distance function to a point cloud [GJ03, CPP08]. For points sampled on a smooth surface, it has been shown that selected stable manifolds provide a suitable reconstruction [DGRS08]. Also, the unstable manifolds have been shown to provide an approximation of the medial axis of the sampled surface [GRS06], with applications to the identification of cylindrical and flat regions [GDB06].

Contributions and paper overview. In spite of their successes, the reconstruction strategies just outlined only accommodate manifold shapes, and since a single reconstruction is reported, are unable to shed light on the potential ambiguities of the sampling. The goal of this paper is to change this paradigm in three ways, by developing an algorithm (i) handling general shapes—compact sets, (ii) reporting a collection of plausible reconstructions, and (iii) selecting the most stable features. Most importantly, the algorithm does not make any assumption on the model being reconstructed, be it its geometry or topology. Moreover, it is effective on challenging datasets, and we believe it is the first such algorithm.

Pre-requisites on the flow complex are recalled in section 2. The algorithm and its theoretical guarantees are respectively exposed in sections 3 and 4. Section 5 illustrates the main features of the algorithm on challenging examples.

2 Background: The Flow Complex and its Hasse Diagram

Morse theory. Morse theory is concerned with the study of functions on manifolds. Following classical terminology in differential topology, a *critical* point of a differentiable function is a point where the differential of the function vanishes, and the function is called a *Morse* function if its critical points are isolated and non-degenerate. The stable (unstable) manifold $W^s(p)$ ($W^u(p)$) of a critical point p is the union of all integral curves associated to the gradient of the function, and respectively ending (originating) at p . The function is termed *Morse-Smale* provided its stable and unstable manifolds intersect transversely [PdM82]. For such a function, the *Morse-Smale complex* is the subdivision of M formed by the connected components of the intersections $W^s(p) \cap W^u(q)$, where p and q range over all critical points. A topological sketch of the manifold can be obtained from a CW complex built from the stable manifolds of the critical points and their incidence.

Morse theory of the distance function to a point cloud. The previous ideas can be instantiated in the following setting. Let P be a finite 3D point set, and denote $d_P(p)$ the distance from any point p to P . This distance function is closely related to the Voronoi diagram of P , which features the points in 3D space equidistant from at least two points in P . It is easily seen that function d_P is smooth everywhere besides at the points in P and on the lower dimensional Voronoi faces. But a generalized gradient ∇d_P can be defined [Lie04], together with the accompanying notions from Morse theory. In particular, a critical point of d_P is a point which is contained in the interior of the convex hull of its nearest neighbors. The stable (unstable) manifolds of these critical points are defined as in the smooth setting. The collection of these stable and unstable manifolds defines the *flow complex* [GJ03].

Of particular interest for our reconstruction algorithm are the stable manifolds. The stable manifold of an index i critical point is i -dimensional¹, and is bounded by the stable manifolds of critical points of index $i - 1$. This recursive structure, which corresponds to incidences between critical points in co-dimension one is encoded in a graph called the *Hasse diagram* of the flow complex: a node associated to an index i critical point a is connected to a node associated to an index $i + 1$ critical point b iff there exists an orbit leaving a and ending at b . In the Hasse diagram, the *successors* (resp. *ancestors*) of node a are denoted $\text{Out}(a)$ (resp. $\text{In}(a)$). See Fig. 1 for a 2D illustration.




Figure 1: Distance function d_P to a 2D point set $P = \{p_0, \dots, p_3\}$. (a) Critical points. (b) Associated Hasse diagram with one node per critical point—the projection of nodes along the x axis is immaterial.

¹For stable manifolds of index two, an additional constraint is required: the stable manifold should not contain any Voronoi vertex.

3 Algorithm

The description of the algorithm is organized as follows: first, we present the building blocks; next, we present the basic version; finally, we present a refined version, which exploits a simplification of the Hasse diagram based upon topological persistence [ELZ02].

3.1 Building Blocks

The algorithm consists of incrementally adding stable manifolds to the reconstruction, a process controlled by a threshold t_r encoding the ratio of critical values of incident critical points. Since these manifolds are selected from the Hasse diagram and since there is a one-to-one correspondence between stable manifolds and nodes of the Hasse diagram, we shall abuse terminology and say that we add *nodes* to the reconstruction. Similarly, the index of a node refers to the index of the critical point associated to this node. Finally, the critical value $V(c)$ associated to a node c is the value of $d_P(c)$ at the corresponding critical point.

If a node c already in the reconstruction triggers the insertion of node d , then node c is called a *sponsor* of node d . As we shall see, a node may have several sponsors, and the discovery of new sponsors occurs through three extension operations called *regularization*, *upflow extension*, *horizontal extension*. The priority associated with a sponsorship relationship is measured by a (regularization / upflow / horizontal) ratio, i.e. a real number ≥ 1 . To describe these operations, we consider two nodes c and d , with node c in the reconstruction.

Reconstruction initialization. The reconstruction is initialized from selected one dimensional stable manifolds. In general, it can be shown that these stable manifolds are exactly the Gabriel edges, namely the Delaunay edges whose diametral ball is empty. In our case, we retain a Gabriel edge $e = (v_0, v_1)$ as an *initialization edge* iff v_1 is the nearest neighbor of v_0 , or vice-versa.

Regularization. The reconstruction shall be a complex. Recall that a stable manifold of an index i critical point is bounded by stable manifolds of index $i - 1$ critical points. Therefore, a node c belonging to the reconstruction is said to be *regularized* if (i) the reconstruction also contains all its ancestors, i.e. all nodes d with $d \in \text{In}(c)$ (ii) these nodes are themselves regularized. A node c is called a *regularization sponsor* of each of its ancestors, and the *regularization ratio* of such a pair (c, d) is set to one, that is $r_r(c, d) = 1$.

Upflow extension. If c and d are incident nodes in the Hasse diagram, with c in the reconstruction, this extension operation consists of inserting node d provided that the ratio between the critical values of the two nodes is bounded by t_r . That is, let c and d be two incident nodes with $\dim(d) = \dim(c) + 1$, and define the *upflow ratio* of the pair as $r_u(c, d) = V(d)/V(c)$. Nodes c and d satisfy the *upflow condition* provided that $r_u(c, d) < t_r$, in which case node c is called an *upflow sponsor* of node d .

Horizontal extension. In this last operation, index one nodes sponsor nodes with the same dimension. More precisely, let c and d be two nodes which are the successors of a common node b in the Hasse diagram, and define the *horizontal ratio* as $r_h(c, d) = V(d)/V(c)$. The pair (c, d) satisfies the *horizontal criterion* provided that $r_h(c, d) < t_r$, in which case node c is termed a *horizontal sponsor* of node d .

3.2 Algorithm Without Persistence

The iterative reconstruction uses a priority queue Q_R which contains all the nodes sponsored as defined above. The priority of a node in the queue is the least upflow / horizontal / regularization ratio over its sponsors included in the reconstruction so far. Nodes in the queue are precisely those which can trigger the insertion of additional nodes thanks to the three extension operations.

More precisely, the queue initialization consists of inserting into Q_R the nodes sponsored by the initialization edges, through upflow and horizontal extension. (The regularization of an edge consists of inserting its vertices into the reconstruction.) Then, the algorithm consists of iteratively popping the node with least priority, so as to perform the regularization, upflow extension, and horizontal extension, in this order. To see how, let c be the node popped. Thanks to the regularization, node c sponsors all its ancestors $\text{In}(c)$. Thanks to the upflow extension, node c sponsors a subset of its successors $\text{Out}(c)$. Thanks to the horizontal extension, node c sponsors a subset of its siblings. Notice that for a pair (c, d) discovered while performing an extension operation, one faces two situations: if node d is already in Q_R , its priority is updated if the ratio of the pair (c, d) is less than that already in the queue; if not, node d is inserted into Q_R with the priority as defined by its sponsor.

The following comments are in order:

- Since the regularization ratio is set to one—the least possible priority, the regularization of a node is immediate and precedes any upflow or horizontal extension.
- The fact that the extension always proceeds with the node associated with the least priority provides a canonical ordering of the nodes found in the reconstruction. In particular, if t_{r_1} and t_{r_2} are two thresholds such that $t_{r_1} < t_{r_2}$, the sequences of operations for the reconstructions associated with t_{r_1} and t_{r_2} are nested.

3.3 Simplifying the Hasse diagram

While the above algorithm is already provably correct, as we will see in the next section, a number of practical situations significantly benefit from a pre-processing aiming at widening the gap between stable manifolds corresponding to the object being reconstructed and those corresponding to its complement. The preprocessing consists of iteratively simplifying the Hasse diagram by a sequence of cancellations, also called *elementary reductions* in [KMM04]. At each iteration, we choose to cancel the pair of incident nodes $e = (a, b)$ with the least ratio $r_u(a, b)$. Since a and b are incident their indices differ by 1. Assume w.l.o.g. that $\dim(b) = \dim(a) + 1$. Consider the bipartite graph over the sets $\text{In}(b)$ and $\text{Out}(a)$. To get the new Hasse diagram, we simply add all edges in the previous graph to the current Hasse diagram, counting multiplicities modulo 2. In particular, nodes a and b are not connected to any node anymore and can thus be removed from the graph. During the cancellation, we also redistribute the stable manifolds attached to the node b to the nodes of $\text{Out}(a)$. The process is illustrated on Fig. 2.




Figure 2: Simplifying the example of Fig. 1. (a) Cancelling the pair (σ_2, M_1) in the Hasse diagram consists of reversing the flow from M_1 to σ_2 . (b) Local structure of the Hasse diagram after the cancellation: dashed edges have been removed. Notice that the stable manifold of the maximum M_1 , which recursively contains the stable manifold of σ_2 and its endpoints has been redistributed to M_2 . (c) Geometrically, reversing the flow from M_1 to σ_2 consists of virtually making the triangle $p_0p_1p_2$ obtuse.

It can be shown that during the simplification process, the pairs are being cancelled in order of increasing topological persistence. We stop the process when the ratio of the next pair exceeds a threshold $t_p \geq 1$. The original algorithm may then be run on the simplified Hasse diagram, retaining for the initialization step only the initial edges that have not been cancelled.

4 Theoretical Guarantees

Our proof of correctness builds on recent inference results obtained in [CCSL06] using the framework of distance functions. Before stating the guarantees of our algorithm, we briefly recall the latter results.

Sampling conditions. Let K be a compact subset of $\mathbb{R}^n$. Though the distance function ² d_K is not differentiable on the medial axis of $\mathbb{R}^3 \setminus K$, it is possible to define a notion of generalized gradient, denoted by ∇d_K , that shares many properties with usual gradients [Lie04]. Denoting by $\Gamma(x)$ the set of points on K closest to $x \in \mathbb{R}^3$, it can be shown that $\|\nabla d_K(x)\|$ is the cosine of the (half) angle of the smallest cone with apex x that contains $\Gamma(x)$. In particular $\|\nabla d_K(x)\|$ equals 1 outside the medial axis of $\mathbb{R}^3 \setminus K$.

²The distance function d_P to a point cloud is a particular case of the distance function d_K to a general compact set.

Definition 1 (μ -reach). *The μ -medial axis of a compact set $K \subset \mathbb{R}^n$ is the set of points $x \notin K$ such that $\|\nabla d_K(x)\| < \mu$. The μ -reach of K , denoted by $r_\mu(K)$, is the minimum distance between a point in K and a point in the closure of its μ -medial axis.*

For $\mu = 1$, the μ -reach coincides with the minimum of the local feature size function [AB99], also called *reach* [Fed59]. The main advantage of the μ -reach over the reach is that it is non-zero—for a suitable value of μ —for a large class of non-smooth shapes, such as polyhedra. Using the concept of μ -reach, one can formulate a sampling condition similar to the ε -sample condition introduced by Amenta et al [AB99]. To state it, we use the following notations. For a positive number α , we denote by K^α the α -offset of K , namely the set of points at distance at most α from K . The Hausdorff distance $d_H(K, P)$ between compact subsets K and P is the least value α such that $K \subset P^\alpha$ and $P \subset K^\alpha$. The condition reads as follows:

Definition 2 ((κ, μ) -approximation). *Given two non-negative real numbers κ and μ , we say that a compact set $P \subset \mathbb{R}^n$ is a (κ, μ) -approximation of a compact set $K \subset \mathbb{R}^n$ if the Hausdorff distance between K and P does not exceed κ times the μ -reach of K .*

Note that polyhedra for example admit finite (κ, μ) -approximation for suitable μ , whereas they do not admit finite ε -samples since they have zero reach. The following theorem, proved in [CCSL06], shows that a compact set can be reconstructed in a topologically correct way from a (κ, μ) -approximation using simple offsets (see [NSW08] for related results in the smooth case):

Theorem 1. *Let $P \subset \mathbb{R}^n$ be a (κ, μ) -approximation of a compact set K . If*

$$\kappa < \frac{\mu^2}{5\mu^2 + 12}$$

then P^α is homotopy equivalent to K^η for sufficiently small η , provided that

$$\frac{4d_H(K, P)}{\mu^2} \leq \alpha < r_\mu(K) - 3d_H(K, P)$$

A key argument in the proof of the above theorem is that the distance function d_P does not have any critical value in the interval $]4d_H(K, P)/\mu^2, r_\mu(K) - 3d_H(K, P)[$. We note that η cannot be set to 0 in general in the theorem due to certain pathological examples.

Correctness. We say that a point cloud P is a ρ -uniform approximation of a compact set K if half the distance between the two closest sample points in P is at least ρ times the Hausdorff distance between P and K , where $0 < \rho < 1$. Using the results from the previous paragraph, the following theorem is proved in section 7.1:

Theorem 2. *Let K be a compact subset of $\mathbb{R}^3$ and assume point cloud P is a ρ -uniform (κ, μ) -approximation of K . If*

$$\frac{4}{\rho\mu^2} < t_r < \frac{\mu^2}{4\kappa} - 1$$

then the output of the algorithm is homotopy equivalent to K^η for small enough η .

It should be noted that the sampling condition in the above theorem is uniform, meaning that it requires that the sampling density is everywhere at least a certain fraction of the global feature size. However, our algorithm itself is adaptive, in the sense that it can cope with situations where no such global density threshold exists.

5 Experiments

This section illustrates key features of our algorithm, whose implementation is sketched in section 7.2.

5.1 Models and Parameters

Models used. We report results on two models: first, two intersecting hemi-spheres (3,000 pts), to illustrate (i) the reconstruction of non-manifold shapes, and (ii) the enumeration of plausible shapes; second, the vase model (2,699 pts), to highlight (i) the importance of persistence to select prominent features, and (ii) the homotopy type preservation. Additional illustrations on a mechanical part (12,593 pts) and a beech tree (20,956 pts) are reported in the supplement in section 7.3. Running times on a standard desktop computer range from 8 seconds for the vase model, to 120 seconds for the beech tree. (Details in section 7.2.)

Parameters. The persistence and reconstruction thresholds t_p and t_r should be finite and larger than one. We thus adopt the following conventions: $t_p = 0$ means that no persistence is used, and $t_r = \infty$ means that all the possible reconstruction steps are carried out. An experimental justification of the values used for t_r and t_p is provided in section 7.4.

Artwork conventions for illustrations. Critical points of index 0/1/2/3 are respectively depicted as grey/yellow/orange/red cubes. Gabriel edges used for the reconstruction initialization are represented as blue line-segments. Index two stable manifolds are depicted as green triangulated surfaces. To represent index three stable manifolds, we display line-segments joining the index three critical point to the index two critical points found on the boundary of the corresponding stable manifold. Regarding persistence, any two points paired by the persistence algorithm are linked by a pink line-segment.

5.2 Results

Accommodating non manifold shapes. As an example of non-manifold reconstruction, consider the two hemi-spheres on Fig. 3(a). As seen from Figs. 3(b,c), the persistence algorithm helps in cancelling maxima that yield a thickening of the reconstruction near the intersection. As seen from Fig. 3(d), the reconstruction of the circle arc found at the intersection consists of a sequence of edges of multiplicity four and cycles mixing edges of multiplicity one, three and five. The former case corresponds to a locally homeomorphic reconstruction. The latter corresponds to a transverse stretching of the intersection circle into a homotopic region, as seen on Fig. 3(e).

Enumerating plausible reconstructions. Reconstruction is an ill-posed problem, which, in general, does not admit a unique solution. As an example, consider Fig. 4(a). Do the circled points feature a hole on the surface or not? Our reconstruction strategy offers the possibility to consider both situations, since increasing t_r results in filling the hole —Fig. 4(b).

Assessing the role of persistence. Models which are noisy and/or under-sampled may feature undesired critical points. The successors and ancestors of these points in the Hasse diagram may yield undesired stable manifolds in the reconstruction.

As an example, the fin of Fig. 5(a) is clearly erroneous in the reconstruction of the vase model. By tracking backward the extension operations performed, as seen from Fig. 6(a,b), it appears that a jump has been made between a surface critical point of index one and an index two critical point located on the fin. On the other hand, this latter critical point is located very close from an index two critical point. Running the persistence algorithm with threshold $t_p = 1.05$ mates the two critical points, and the jump from the surface critical point is now beyond the threshold $t_r = 1.7$, as seen from Fig. 5(b). The effect of persistence is also easily seen from the so-called reconstruction profiles, presented in the supplemental-section 7.4.

Notice however, that persistence may prevent from inserting selected stable manifolds in the reconstruction. To see why, consider an index two critical point a paired by the persistence algorithm to a maximum M , and also assume that the outflow of a consists of M and of the critical point at infinity. When updating the Hasse diagram as explained in section 3.3, the stable manifold of b is associated to nodes in the outflow of node a , that is to the maximum at infinity. Thus, the stable manifold of node a cannot be part of the reconstruction since the maximum at infinity is never reached. Should this happen, notice that the size of the hole on the surface is comparable to the size of the 3D hole associated to the index 3 critical point paired, a characterization of under-sampling.

Finally, notice that the possibility to track back the extension operations, which is illustrated on Fig. 6(a,b), offers a unique way to control under-sampled regions in the point cloud processed, so as to further guide the acquisition process.




Figure 3: Reconstructing non-manifold shapes (a) Reconstruction of two intersecting hemi-spheres with $t_r = 1.9$, $t_p = 0$. (b) Transparent view of a section of the reconstruction in (a): the stable manifolds of the circled maxima correspond to a thickening of the reconstruction. (c) Same region as in (b) with parameters $t_r = 2.5$, $t_p = 1.05$: the maxima have been cancelled by the persistence algorithm. (d) Reconstruction for parameters $t_r = 2.5$, $t_p = 1.05$: Gabriel edges of multiplicity zero (green), one (purple), three (blue), four (orange) and five (yellow). (e) Circled region of Fig. (d): the intersection curve from p_0 to p_1 has been stretched to a topological disk, namely the union of the stable manifolds of the index two critical points c and d —triangles of the stable manifolds not shown.



Figure 4: Enumerating plausible reconstructions (a) At $t_r = 1.9$, $t_p = 0$, the red-circled points punch a hole into the surface of Fig. 3(a). (b) At $t_r = 2.1$, the hole has been filled by the stable manifolds of the index two critical points c and d .



Figure 5: Assessing the importance of persistence (a) At $t_r = 1.7$, $t_p = 0$, a fin between two handles of the vase model is observed, while one back handle is disconnected—circled region. (b) At $t_r = 2$, $t_p = 1.02$, the fin is gone. Note also the preservation of the homotopy type of a solid handle, which is reconstructed as a polyline.




Figure 6: Untangling the role of persistence in Fig. 5(a,b). (a,b) The extension path followed by the algorithm—red arrows. The circled region features an upflow extension between a surface critical point c and a medial axis critical point d . (c) Using persistence, the critical point d has been paired to an index one critical point e . The upflow extension from the surface critical point c is now forbidden at threshold $t_r = 2$, whence the reconstruction of Fig. 5(b).

6 Discussion and Outlook

This paper presents a novel reconstruction algorithm, characterized by the following features: (i) The algorithm accommodates the reconstruction of non-manifold shapes, and does not rely on any concept from the smooth setting. It only relies on properties of the distance function to samples. (ii) Under mild hypothesis, it is proved to provide correct reconstructions in terms of homotopy type. (iii) It does not provide a single reconstruction, but instead allows one to enumerate all *plausible* reconstructions, with respect to a threshold encoding the *proximity* between critical points of the distance function. (iv) The algorithm allows the selection of prominent features thanks to a simplification of the Hasse diagram based upon persistence. (v) In case of reconstruction *failure*, the user can get insights in the structure of the sampling by tracking back the sequence of reconstruction extensions, so as to locate the under-sampled area responsible for the failure and possibly fix it.

In spite of these novel features, a number of questions deserve further work. Complexity-wise, our algorithm relies on the flow complex, whose construction is the bottleneck. However, one could design a strategy interleaving the flow complex construction and the reconstruction, so as to compute the former on demand. Such an algorithm would have output-sensitive complexity, since *one step ahead* only would be required with respect to the final reconstruction.

In terms of output, the reconstruction consists of stable manifolds of the flow complex. In particular, the surface patches correspond to index two stable manifolds, whose triangles may have arbitrary aspect ratio. This observation calls for further work so as to approximate the surface patches of the reconstruction using Delaunay triangles.

Finally, since our algorithm accommodates non manifold shapes, it should allow the development of strategies computing stratifications of complex shapes, with applications to morphological analysis.

References

- [AB99] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete Comput. Geom.*, 22(4):481–504, 1999.
- [ABCO⁺01] M. Alexa, J. Behr, D. Cohen-Or, S. Fleishman, D. Levin, and C. Silva. Point set surfaces. In IEEE, editor, *Visualization*, 2001.
- [ACDL00] N. Amenta, S. Choi, T. K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. In *Proc. 16th Annu. Sympos. Comput. Geom.*, pages 213–222, 2000.
- [BC02] J.-D. Boissonnat and F. Cazals. Smooth surface reconstruction via natural neighbour interpolation of distance functions. *Comp. Geometry Theory and Applications*, 22(1):185–203, 2002. Conf. version: ACM Sympos. Comput. Geom.'00.
- [Boi84] J-D. Boissonnat. Geometric structures for three-dimensional shape representation. *ACM Trans. Graph.*, 3(4):266–286, 1984.
- [CBC⁺01] J. Carr, R. Beatson, J. Cherrie, T. Mitchell, T. Fright, B. McCallum, and T. Evans. Reconstruction and representation of 3d objects with radial basis functions. In *Siggraph*. ACM, 2001.

- [CCSL06] F. Chazal, D. Cohen-Steiner, and A. Lieutier. A sampling theory for compacts in Euclidean spaces. In *ACM Symp. Comp. Geometry*, pages 319–326, 2006.
- [CDA⁺08] J.-C. Chambelland, M. Dassot, B. Adam, N. Donès, P. Balandier, A. Marquier, M. Saudreau, G. Sonohat, and H. Sinoquet. A double-digitising method for building 3d virtual trees with non-planar leaves: application to the morphology and light-capture properties of young beech trees (*fagus sylvatica*). *Functional Plant Biology*, 35(9 & 10):1059–1069, Novembre 2008.
- [CG06] F. Cazals and J. Giesen. Delaunay triangulation based surface reconstruction. In J.-D. Boissonnat and M. Teillaud, editors, *Effective Computational Geometry for curves and surfaces*. Springer-Verlag, Mathematics and Visualization, 2006.
- [Cha03] R. Chaine. A convection geometric-based approach to surface reconstruction. In *Symp. Geometry Processing*, pages 218–229, 2003.
- [CPP08] F. Cazals, A. Parameswaran, and S. Pion. Robust construction of the three-dimensional flow complex. In *ACM Symposium on Computational Geometry*, pages 182–191, 2008.
- [CSD04] D. Cohen-Steiner and F. Da. A greedy delaunay based surface reconstruction algorithm. *The Visual Computer*, 20(1):4–16, 2004.
- [DGJ03] T. Dey, J. Giesen, and M. John. Alpha-shapes and flow shapes are homotopy equivalent. In *ACM Symposium on Theory of Computing*, pages 493—501, 2003.
- [DGRS08] T.K. Dey, J. Giesen, E. Ramos, and B. Sadri. Critical points of the distance to an epsilon-sampling of a surface and flow based surface reconstruction. *Int'l Journal of Computational Geometry and Applications*, 18:29–61, 2008.
- [Ede95] H. Edelsbrunner. The union of balls and its dual shape. *Discrete & Computational Geometry*, 13:415–440, 1995.
- [Ede04] H. Edelsbrunner. Surface reconstruction by wrapping finite point sets in space. *Discrete and Computational Geometry*, 32:231–244, 2004.
- [ELZ02] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete and Computational Geometry*, 28:511–533, 2002.
- [Fed59] H. Federer. Curvature measures. *Transactions of the American Mathematical Society*, 93(3):418–491, 1959.
- [GDB06] S. Goswami, T. K. Dey, and C. Bajaj. Identifying flat and tubular regions of a shape by unstable manifolds. In *ACM Solid and Physical Modeling*, 2006.
- [GJ02] J. Giesen and M. John. Surface reconstruction based on a dynamical system. In *Proceedings of the 23rd Annual Conference of the European Association for Computer Graphics (Eurographics), Computer Graphics Forum 21*, pages 363–371, 2002.
- [GJ03] J. Giesen and M. John. The flow complex: A data structure for geometric modeling. In *ACM SODA*, 2003.
- [GRS06] J. Giesen, E. Ramos, and B. Sadri. Medial axis approximation and unstable manifolds. In *ACM SoCG*, 2006.
- [KMM04] T. Kaczynski, K. Mischaikow, and M. Mrozek. *Computational Homology*. Springer, 2004.
- [Lie04] A. Lieutier. Any open bounded subset of $\mathbb{R}^n$ has the same homotopy type than its medial axis. *Computer-Aided Design*, 36(11):1029–1046, 2004.
- [NSW08] P. Niyogi, S. Smale, and S. Weinberger. Finding the homology of submanifolds with high confidence from random samples. *Discrete & Computational Geometry*, 39(1-3):419–441, 2008.
- [PdM82] J. Palis and W. de Melo. *Geometric Theory of Dynamical Systems*. Springer, 1982.
- [ZOF01] H. K. Zhao, S. Osher, and R. Fedkiw. Fast surface reconstruction using the level set method. In *Proceedings of IEEE Workshop on Variational and Level Set Methods in Computer Vision (VLSM)*, 2001.

7 Supplemental

7.1 Proof of theorem 2

We prove the theorem for the basic algorithm, without resorting to Hasse diagram simplification. It is apparent from the proof that the guarantee also holds for the full version of the algorithm, as long as the chosen persistence threshold t_p does not exceed t_r .

Proof. Since ρ and μ do not exceed 1, the condition implies that $\mu^2/(4\kappa) - 1 > 4$, so $\kappa < \mu^2/20$. As $\mu \leq 1$, $20 > 5\mu^2 + 12$, so the reconstruction theorem can be applied. In particular, d_P does not have any critical value in the interval $]4d_H(K, P)/\mu^2, r_\mu(K) - 3d_H(K, P)[$.

We now show the output of our reconstruction algorithm is the union of the stable manifolds of critical points of d_P whose distance to P does not exceed $4d_H(K, P)/\mu^2$. Let c be such a critical point. Consider a downward path starting at c and ending at an index 1 critical point d in the Hasse diagram of P 's flow complex. We clearly have $d_P(d) < d_P(c)$. Let x be one of the sample points in the boundary of the stable manifold of d , and e be the edge joining x and its closest sample point. Edge e belongs to the reconstruction by construction. Also, the critical value corresponding to e (i.e. its half length) is at least $\rho d_H(K, P)$ by assumption. Hence the ratio between the (half) lengths of e and of the stable manifold of d is at most $4/(\rho\mu^2) < t_r$, which shows that the stable manifold of d must have been included in the reconstruction through a horizontal extension step. Now the ratio $d_P(c)/d_P(d)$ is also at most $4/(\rho\mu^2)$, so the stable manifold of c must have been included in the reconstruction through upflow extensions from d .

It is not difficult to show that stable manifolds of critical points with value at least $r_\mu(K) - 3d_H(K, P)$ do not belong to the reconstruction. Indeed, this would imply an upflow extension from a critical point with value at most $4d_H(K, P)/\mu^2$, but the ratio between the two values is:

$$\begin{aligned} \frac{r_\mu(K) - 3d_H(K, P)}{4d_H(K, P)/\mu^2} &= \frac{\mu^2}{4} \left(\frac{r_\mu(K)}{d_H(K, P)} - 3 \right) \\ &\geq \frac{\mu^2}{4} (\kappa^{-1} - 3) \\ &\geq \frac{\mu^2}{4\kappa} - 1 \end{aligned}$$

which is larger than threshold t_r by assumption. To conclude the proof of the theorem, it is sufficient to use the fact that the union of the stable manifolds of critical points with value less than some threshold α is homotopy equivalent to the α -shape of the point cloud P [DGJ03], or equivalently to its α -offset P^α [Ede95]. Indeed, from the reconstruction theorem, the latter offset is homotopy equivalent to K^η for sufficiently small η . $\square$

7.2 Implementation Outline

Our implementation of the reconstruction algorithm meets the C++ standards of Computational Geometry Algorithms Library (CGAL, see www.cgal.org) and is parametrized by a traits class corresponding the Hasse diagram of the distance function to the samples. The construction of this Hasse diagram is carried out with our implementation of the flow complex [CPP08]. More precisely, incidences between index one and index two critical points are discovered while building the stable manifolds of the latter; incidences between index two and index three critical points are discovered while computing the unstable manifolds of the former. As reported in [CPP08], this latter stage is the limiting step, due to cascaded constructions and predicates on such.

The code was compiled with CGAL 3.3, using GMP 4.2, with the GNU compiler g++ 4.1.2. The machine used for the experiments was a PC running Linux Fedora Core 7, with 2MB of memory and a 3GHz Pentium 4 processor.

7.3 Additional Illustrations

This section provides illustrations on two additional models: the mechanical part of Fig. 7 is a typical model in computer aided geometric design / reverse engineering; the beech tree of Fig. 8 is an example of a scanned plant in agricultural sciences [CDA⁺08].




Figure 7: Reconstruction of a mechanical part—a genus 3 surface, with parameters $t_r = 2.1, t_p = 1.1$. Inset: circled points make up the boundary (a topological circle) of the cylindrical depression to the right of the largest hole.




Figure 8: Beech tree reconstruction, with parameters $t_r = 2.2, t_p = 1.1$ (a) Overview of this noisy and under-sampled model (b,c) Zoom near a an under-sampled peduncle. The point cloud is courtesy of J-C. Chambelland et al, UMR 547 PIAF - INRA/UBP.

7.4 Reconstruction Profiles

As observed in section 3.2, the reconstructions for two thresholds t_{r_1} and t_{r_2} such that $t_{r_1} < t_{r_2}$ are nested. Assume that at the i th iteration on Q_R , a node with priority r_i is popped. The sequence $(r_i)_{i \geq 1}$ for $t_r = \infty$ encompasses all possible reconstructions, and is termed the *reconstruction profile* of the model. The interests of profiles are twofold.

First, profiles provide an easy assessment of the influence of the simplification of the Hasse diagram. When no persistence is used, situations where one extension with arbitrary extension ratio is followed by a number of extensions with ratio close to one are frequent: these situations correspond to the numerous crenels observed

on Figs. 9(a), 11(a), 13(a), 15(a). When persistence is used, since critical points whose critical values are close get matched, these crenels disappear even with a modest persistence threshold—Figs. 9(b), 11(b), 13(b), 15(b).

Second, profiles encode the magnitude of the cuts between the critical points located on/in the objects and in its complement. As seen from Fig. 10 for the two hemi-spheres, one does not see any significant gap between both sets of critical points. The distance function does not have any critical point on the medial axis of this model. The situation is different on the vase and the mechanical part, as seen from Figs. 12 and 14. For these models, one indeeds observes a clear cut between the critical values, which accounts for the reconstruction threshold taken around $t_r = 2$. Notice in particular that in both cases, the number of critical points on/in the object and in its complement are *incommensurable*: while the former are related to the sampling density, the latter are related to the number *features* of the complement. Finally, the beech tree model—see Fig. 16 corresponds to a more complex setting, where several *reconstruction scales* coexist.

In passing, we note that the quality of the reconstruction is not sensitive to the thresholds used: values of $t_p \in [1.02, 1.2]$ and $t_r \in [1.9, 2.1]$ yield comparable results.

7.4.1 Reconstruction Profiles: Intersecting Spheres



Figure 9: Reconstruction profiles for intersecting spheres (a) $t_p = 0, t_r = \infty$ (b) $t_p = 1.1, t_r = \infty$




Figure 10: Zoom of Fig. 9(b)

7.4.2 Reconstruction Profiles: Vase



Figure 11: Reconstruction profiles for the vase (a) $t_p = 0, t_r = \infty$ (b) $t_p = 1.1, t_r = \infty$




Figure 12: Zoom of Fig. 11(b)

7.4.3 Reconstruction Profiles: Mechanical part



Figure 13: Reconstruction profiles for the mechanical part (a) $t_p = 0, t_r = \infty$ (b) $t_p = 1.1, t_r = \infty$




Figure 14: Zoom of Fig. 13(b)

7.4.4 Reconstruction Profiles: Beech Tree



Figure 15: Reconstruction profiles Beech tree (a) $t_p = 0, t_r = \infty$ (b) $t_p = 1.1, t_r = \infty$




Figure 16: Zoom of Fig. 15(b)

Contents

1	Introduction	3
2	Background: The Flow Complex and its Hasse Diagram	4
3	Algorithm	5
3.1	Building Blocks	5
3.2	Algorithm Without Persistence	5
3.3	Simplifying the Hasse diagram	6
4	Theoretical Guarantees	6
5	Experiments	7
5.1	Models and Parameters	8
5.2	Results	8
6	Discussion and Outlook	10
7	Supplemental	12
7.1	Proof of theorem 2	12
7.2	Implementation Outline	12
7.3	Additional Illustrations	12
7.4	Reconstruction Profiles	13
7.4.1	Reconstruction Profiles: Intersecting Spheres	15
7.4.2	Reconstruction Profiles: Vase	15
7.4.3	Reconstruction Profiles: Mechanical part	16
7.4.4	Reconstruction Profiles: Beech Tree	16



Unité de recherche INRIA Sophia Antipolis
2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Unité de recherche INRIA Futurs : Parc Club Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 ORSAY Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)

Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399