

Sampling and Reconstructing Manifolds Using Alpha-Shapes

Fausto Bernardini Chandrajit L. Bajaj

Department of Computer Sciences
Purdue University*

Abstract

There is a growing interest for the problem of reconstructing the shape of an object from multiple range images. Several methods, based on heuristics, have been described in the literature. We propose the use of alpha-shapes, which allow us to give a formal characterization of the reconstruction problem and to prove that, when certain sampling requirements are satisfied, the reconstructed alpha-shape is homeomorphic to the original object and approximate it within a fixed error bound.

In a companion paper, we describe practical methods to automatically select an optimal alpha value, to deal with less-than-ideal scans, and to fit smooth piecewise algebraic surface to the data points.

1 Introduction

Cheaper, easier-to-use 3D digitizers are fostering a growing interest for the problem of *shape-reconstruction*. Automatic methods for reconstructing an accurate geometric model of an object from a set of digital scans have applications in reverse engineering, shape analysis, virtual worlds authoring, 3D faxing and tailor-fit modeling.

Range or optical-triangulation laser scanners produce a regular grid of measurements, which can be easily converted to a rectangular or cylindrical surface model when a single scan suffices to capture the whole object's surface. However more often multiple scans are required, and the results must be merged together. Several approaches have been proposed to reconstruct the shape of an object from a collection of digital scans.

Turk and Levoy [14] proposed to “zipper” together several meshes obtained from separate 3D-scans of an object. More recently, Curless and Levoy [5] presented an approach to merge several range images by scan-converting each image to a weighted signed distance function in a regular 3D grid. The zero-contour of the signed distance function, which can be easily extracted with a marching cubes algorithm [13], represents the reconstructed surface.

A different class of methods try to rely on spatial location of points only, without any assumed knowledge of connectivity between sampled points. Boissonnat [3] proposes two methods to build a triangulation having the given points as vertices. Following his first approach, one starts with creating an edge between the two closest points. A third point is then chosen and added, so that a triangle is formed. Other points are successively added and new triangles are created, and joined to an edge of the current triangulation boundary, until all points have been included. The second method is based on the idea of first computing a Delaunay triangulation of the convex hull of the set of points, and then *sculpturing* the volume by removing tetrahedra, until all points are on its boundary, or no tetrahedra can be further removed.

Choi *et al.* [4], described a method to incrementally form a triangulation interpolating all data points, based on the assumption that there exists a point from which all the surface is visible. After a triangulation is built, it is improved by edge swapping based on a smoothness criterion.

Veltkamp [15] introduced a new general geometric structure, called γ -graph. The γ -graph coincides initially with the convex hull of the data points, and is progressively *constricted* (i.e. tetrahedra having boundary faces are deleted) until the boundary of the γ -graph is a closed surface, passing through all the given points.

Hoppe *et al.* [11] compute a signed distance function from the data points, and then use its zero-

*West Lafayette, Indiana
47907-1398 USA

{fxb,bajaj}@cs.purdue.edu
<http://www.cs.purdue.edu/research/shastra>

contour as an approximation of the object. To define the signed distance from the unknown surface, they compute a best-fit tangent plane for each data point, and then find a coherent orientation for the surface by propagating the normal direction from point to point, using a precomputed minimum spanning tree to favor propagation across points whose associated normals are nearly parallel.

One of the most difficult problems of shape reconstruction from unorganized points is understanding how to “connect-the-dots” so as to form a surface that has the same topological (e.g. number of handles) and geometric (e.g. depressions and protrusions) characteristics of the original. All the methods listed above are based on geometric heuristics. While these methods have been shown to be successful on several examples and practical applications, they fail to provide requirements on the sampling that guarantee a provably correct reconstruction.

Alpha-shapes were introduced in the plane by Edelsbrunner *et al.* in [8] and then extended to higher dimensions [7, 9], as a geometric tool for reasoning about the “shape” of an unorganized set of points. They offer the dual benefit of having a solid mathematical foundation and of being relatively easy to compute. We have developed several automatic reconstruction methods based on alpha-shapes and algebraic-patch fitting [1, 2].

In this paper we formalize the shape reconstruction problem, give a set of sufficient conditions for reconstructing an object using alpha-shapes, and discuss some practical considerations. The paper is organized as follows: Section 2 contains a short review of the main concepts and notation used in this work. In Section 3 we give a formal statement of the shape reconstruction problems. Section 4 is devoted to a proof of sufficient conditions on the sampling to allow a homeomorphic, error-bounded reconstruction. In section 5 we illustrate some examples, discuss some practical considerations, and outline directions for future work.

2 Preliminaries

Topological spaces, homeomorphisms, and manifolds. A *topological space* is a set S together with a collection \mathcal{U} of subsets of S (that is, \mathcal{U} is a subset of 2^S) satisfying the following conditions:

1. $\emptyset \in \mathcal{U}, S \in \mathcal{U}$.
2. If $U_1, \dots, U_n \in \mathcal{U}$ then $\bigcap_{i=1}^n U_i \in \mathcal{U}$.
3. Arbitrary unions of elements in \mathcal{U} lie in \mathcal{U} ; that is, if $\tilde{\mathcal{U}} \subset \mathcal{U}$, then $\bigcup_{U \in \tilde{\mathcal{U}}} U \in \mathcal{U}$.

The elements of \mathcal{U} are called *open sets* in S . The collection \mathcal{U} is called a *topology* on S . We often suppress the \mathcal{U} and simply refer to S as a topological space.

A map f from a topological space X to another topological space Y is *continuous* if every neighborhood of $f(p)$ in Y is mapped by f^{-1} to a neighborhood of p in X . If f is bijective, and if both f and f^{-1} are continuous, then f is a *homeomorphism*. Two topological spaces X and Y are *homeomorphic* if there exists a homeomorphism $f : X \rightarrow Y$.

In the following, we will restrict ourselves to subsets of the n -dimensional Euclidean space, $S \subset \mathbf{R}^n$. Let us define the following subspaces of \mathbf{R}^n , with origin o :

$$\begin{aligned} H^n &= \{x \in \mathbf{R}^n \mid x_n \geq 0\} \\ B^n &= \{x \in \mathbf{R}^n \mid \|x - o\| \leq 1\} \\ S^{n-1} &= \{x \in \mathbf{R}^n \mid \|x - o\| = 1\} \end{aligned}$$

Open, half-open, and closed n -balls are homeomorphic to \mathbf{R}^n , H^n and B^n , respectively. An $(n-1)$ -sphere is homeomorphic to S^{n-1} .

A set in \mathbf{R}^n is *bounded* if it is contained in an open ball. An *open covering* of a topological space S is a collection $\mathcal{V} \subset \mathcal{U}$ such that $\bigcup_{V \in \mathcal{V}} V = S$. A space S is *compact* if every open covering has a finite sub-covering. A subspace of \mathbf{R}^n that is both closed and bounded is compact.

A k -manifold in \mathbf{R}^n ($n \geq k$) is a subspace that is locally homeomorphic to \mathbf{R}^k . A k -manifold with boundary is a subspace that is locally homeomorphic to either \mathbf{R}^k or the half-open k -ball H^k . Points with a neighborhood homeomorphic to H^k form the *boundary* of the manifold X , denoted $\text{bd}(X)$. The boundary of a k -manifold with boundary is a $(k-1)$ -manifold without boundary.

Simplicial complexes. A k -simplex $\sigma_T = \text{conv}(T)$ is the convex combination of an affinely independent point set $T \subset \mathbf{R}^n$, $|T| = k + 1$. k is the *dimension* of simplex σ_T . A (geometric) *simplicial complex* K is a finite collection of simplices with the following two properties:

1. if $\sigma_T \in K$ then $\sigma_U \in K, \forall U \subset T$
2. if $\sigma_U, \sigma_V \in K$, then $\sigma_{U \cap V} = \sigma_U \cap \sigma_V$ (1 and 2 imply that $\sigma_{U \cap V} \in K$).

The underlying space of K is $[K] = \bigcup_{\sigma \in K} \sigma$. A subcomplex of K is a simplicial complex $L \subset K$.

Alpha-shapes. Alpha-shapes [8, 9] associate a mathematically defined meaning to the vague concept of *shape* of an unorganized set of points. Weighted alpha-shapes [7] are a generalization of alpha-shapes to sets of weighted points. In the following, we will shortly review definitions and properties of alpha-shapes. The presentation is adapted from [7]. Notice that although the exposition is for *unweighted* alpha-shapes, we will use the notation used in the more general weighted case. A weighted alpha-shape coincides with an unweighted alpha-shape when all weights are equal to zero. We restrict our presentation to the three-dimensional case. n -dimensional weighted alpha-shapes are described in the cited reference [7].

In the following we will sometimes regard a sphere of radius ρ centered in p as a weighted point p of weight $w_p = \rho^2$. We define the power distance of a point x from a weighted point p as

$$\pi_p(x) = \|p - x\|^2 - w_p$$

where $\|p - x\|$ is the Euclidean distance between p and x . A geometric interpretation of the power distance is the following: If weighted point p represents a sphere of center p and radius $\sqrt{w_p}$, then $\pi_p(x)$ is the square of the length of a tangent line segment from x to the sphere (see Figure 1).




FIGURE 1: Power distance of a point x from the weighted point p .

Let $P \subset \mathbf{R}^3$ be a finite set of points (general position is assumed implicitly throughout the paper), $|P| \geq 4$, and \mathcal{T} its Delaunay triangulation. For every simplex $\sigma_T \in \mathcal{T}$, let y_T be the smallest sphere (weighted point) such that $\pi_{y_T}(p) = 0, \forall p \in T$. If $|T| = 4$ there is only one such sphere y_T , the circumsphere of σ_T . If $|T| = k + 1 < 4$ there are infinitely many such spheres, but only one has minimum radius. The center of y_T is located at the intersection of the *chordale* of T (see Figure 2)




FIGURE 2: The collection of spheres containing the two vertices of the 1-simplex $T = \{p, q\}$. The sphere y_T of minimum radius ρ_T is drawn in bold.

$$\chi_T = \bigcap_{p, q \in T} \chi_{p, q}, \quad \chi_{p, q} = \{x \in \mathbf{R}^3 \mid \|p - x\| = \|q - x\|\}$$

with the orthogonal k -flat $\text{aff}(T)$. Let ρ_T be the radius of y_T , and call $w_{y_T} = \rho_T^2$ the *size* of the k -simplex σ_T . Notice that the size of a 0-simplex is 0. The size of simplices satisfies the following monotonicity property: if $U \subset T$ then $w_{y_U} < w_{y_T}$, that is the size of a proper face of a simplex is smaller than the size of the simplex itself.

A point $q \in P - T$ is a *conflict* for y_T if $\pi_{y_T}(q) < 0$, and y_T is *conflict-free* if it has no conflicts. Obviously, all 3-simplices $\sigma_T \in \mathcal{T}$ are conflict-free, but a k -simplex, $k < 3$, can have conflicts.

Definition 2.1 *The alpha-complex of P is the subcomplex Σ_α of \mathcal{T} formed by all simplices σ_T such that:*

- (a) *The size of y_T is less than α and y_T is conflict-free, or*
- (b) *σ_T is a face of σ_U and $\sigma_U \in \Sigma_\alpha$.*

The underlying space \mathcal{S}_α of Σ_α , called alpha-shape, is a polytope, which can be non-connected and different from the closure of its interior (i.e. it may contain parts of heterogeneous dimensionality).

It can be proved (see [7]) that the following is an alternative definition of alpha-shapes:




FIGURE 3: An example of ambiguous 2D reconstruction from points. From left to right: A point sampling and three, equally acceptable, reconstructions.

Definition 2.2 Consider a subset $T \subseteq P$, with $|T| = k + 1 \leq 3$, and the k -simplex σ_T . Let us call σ_T α -exposed if there exists a weighted point x , of weight $w_x = \alpha$ (that is, a sphere of radius $\sqrt{\alpha}$), such that

$$\pi_x(p) = \begin{cases} = 0 & \forall p \in T \\ > 0 & \forall p \in P - T \end{cases}$$

The alpha-shape \mathcal{S}_α of P is a polytope whose boundary is the union of all α -exposed simplices spanned by subsets $T \subseteq P, |T| \leq 3$. The interior of \mathcal{S}_α is formed by those components of \mathbf{R}^3 bounded by collections of α -exposed 2-simplices σ_T , such that σ_T is α -exposed only on one side (i.e. there exists only one weighted point of weight α that exposes σ_T). The interior points of \mathcal{S}_α lie on the side of σ_T that is not α -exposed.

When the alpha-shape of a point set P , for some α , is a connected 2-manifold without boundary, it partitions the space into two connected components. The bounded component is the *solid* enclosed by the alpha-shape, and will be called *alpha-solid*. In [2] we give a more general definition of alpha-solid, as well as techniques to automatically find an optimal α value for a given set of points, and heuristics to improve the alpha-solid in areas of insufficient sampling density (see also Section 5).

3 Statement of the Problem

Reconstructing the shape of an object from an unorganized “cloud” of points is in general an under-constrained problem. Consider the simple 2D reconstruction problem illustrated in Figure 3: Several solutions are possible, and it is difficult to identify the “best” among them. It is therefore of interest looking at the following problem: What are the characteristics of a sampling S (a finite set of points) of the surface of a solid object M , such that M can

be reconstructed from S unambiguously and within predefined approximation bounds?

In particular, we consider the following *reconstruction problem*: Starting with a sampling of the surface B of a solid, we want to compute a triangulated surface K that has the “same shape” of B , and such that a suitably defined *distance* $D(K, B)$ of K from B is bounded by a given ε . A useful distance measure is for example:

$$D(K, B) = \max_{p \in [K]} \min_{q \in B} \|p - q\|.$$

Stated formally:

Problem 3.1 Let B be a compact 2-manifold without boundary (in particular, the boundary of a solid M), and $S \subset B$ a finite set of points (sampling). Construct a (geometric) simplicial complex K , such that $K^{(0)} = S$, K is homeomorphic to B , and $D(K, B) < \varepsilon$, for a fixed $\varepsilon \in \mathbf{R}, \varepsilon > 0$.

The pair (K, h) , where K is a simplicial complex and h is a homeomorphism $h : [K] \rightarrow B$ is called a *triangulation* in algebraic topology.

An algorithm aimed at reconstructing the shape of an object from point data alone must have a way of inferring spatial relationships among points. Characteristics of the sampling that guarantee an unambiguous and correct reconstruction depend on how the data is interpreted by the algorithm.

We have already mentioned that alpha-shapes allow us to find spatial relationships between points of an unorganized set. The relationships are based on proximity. Clusters of points close to each other are grouped to form edges, triangles and tetrahedra, and more complex structures made of collections of these simple constituents.

The question we need to answer is therefore the following: What are sufficient conditions of a sampling that guarantee that there exists an α such that the corresponding α -shape satisfies the requirements of Problem 3.1?




FIGURE 4: Sampling requirements for 1-manifolds in \mathbf{R}^2 . (a) The sampling density must be such that the center of the “disk probe” is not allowed to cross B without touching a sample point. (b) The radius ρ of the disk probe must be small enough that the intersection with B has at most one connected component. (c) Examples of non admissible cases of probe-manifold intersections.

4 Sampling Requirements

We can look at the two-dimensional case to get some insight into the problem. Figure 4 illustrates the discussion that follows. In this case, we are sampling a 1-manifold B (observe that B is a collection of “loops”). Intuitively, we can think of the points of the sampling as “pins” that we fix on B . We now use a disk probe of radius $\rho = \sqrt{\alpha}$ to “sense” the manifold. The probe must be able to move from point to point of S on the surface, touching pairs of points in sequence, and without touching other parts of B . The pairs of points will be connected by segments of the alpha-shape, and will form loops homeomorphic (and geometrically close) to each component of B .

Clearly, a necessary condition is that no two adjacent points of the sampling are farther away than the diameter of our disk-probe, because otherwise the probe would “fall” inside the boundary of our solid object. We also need to make sure that all, and only, the edges connecting pairs of adjacent points are α -exposed. To do this, our probe needs to be small enough to be able to isolate a neighborhood of a point p on B , or, equivalently, discern “adjacent” points on B from points that are close in the Euclidean sense but not on the surface. These requirements are formalized in the following

Theorem 4.1 *Let $B \subset \mathbf{R}^2$ be a compact 1-manifold without boundary, and $S \subset B$ a finite point set. If*

1. *For any closed disk $D_\rho \subset \mathbf{R}^2$ of radius ρ , $B \cap D_\rho$ is either (a) empty; (b) a single point p (then $p \in \text{bd}(D_\rho)$); (c) homeomorphic to a closed 1-ball I , such that $\text{int}(D_\rho) \cap B = \text{int}(I)$;*

2. *An open disk of radius ρ centered on B contains at least one point of S ,*

then the alpha-shape \mathcal{S}_α of S , $\alpha = \rho^2$, is homeomorphic to B and

$$D(\mathcal{S}_\alpha, B) = \max_{p \in \mathcal{S}_\alpha} \min_{q \in B} \|p - q\| < \rho.$$

Observe that B is in general a collection of 1-spheres B_i . We will prove the theorem by showing that for each 1-sphere $B_i \subset B$ there is a homeomorphic component in \mathcal{S}_α , and then showing the bound on the distance.

Before we prove the theorem, we need a few lemmas. In the lemmas that follow, B , B_i , and S are those defined above. The symbol D_ρ is used as above to indicate a closed disk of radius ρ . We refer to the two conditions stated in the theorem as conditions 1 and 2. We often refer to two points p, q on a component B_i of B , and use the symbols X, Y to indicate the two closed 1-balls on the 1-sphere B_i having p, q as boundary points. Obviously $X \cup Y = B_i$.

Lemma 4.1 *Let p, q be two points on B . If there exists D_ρ such that $p, q \in \text{bd}(D_\rho)$, then $D_\rho \cap B$ is a 1-ball I , and $\text{bd}(I) = \{p, q\}$.*

Proof: Since D_ρ contains two points of B , by condition 1 it must intersect B in a (closed) 1-ball I , with $p, q \in I$. Suppose $p \notin \text{bd}(I)$. Then $p \in \text{int}(I)$. But $p \notin \text{int}(D_\rho) \cap B$, therefore condition 1 cannot be satisfied. \diamond

Lemma 4.2 *Let p, q be two points on B_i , and let X, Y be the two 1-balls on B_i , $\text{bd}(X) = \text{bd}(Y) =$*

$\{p, q\}$. If there exists D_ρ such that $p, q \in \text{bd}(D_\rho)$, then either $D_\rho \cap B = X$ or $D_\rho \cap B = Y$.

Proof: By Lemma 4.1, $D_\rho \cap B$ is a 1-ball whose boundary is $\{p, q\}$. Clearly this 1-ball must be a subset of B_i . There are only two 1-balls on B_i having $\{p, q\}$ as boundary, namely X and Y . \diamond

Lemma 4.3 *Let p, q be two points on B_i , and let X, Y be the two 1-balls on B_i , $\text{bd}(X) = \text{bd}(Y) = \{p, q\}$. If $\text{int}(X) \cap S = \emptyset$ then $\|p - q\| < 2\rho$.*

Proof: Suppose that $\|p - q\| \geq 2\rho$. Since X is a 1-ball connecting p and q and $\|p - q\| \geq 2\rho$, there exists a point $c \in X$ such that $\|p - c\| = \rho$. Consider D_ρ centered in c , and observe that $p \in \text{bd}(D_\rho)$. Since D_ρ contains two points of B (p and c), it must intersect B in a 1-ball I , and p must be a boundary point of I , by condition 1.

The other boundary point of I must be contained in the 1-ball Z between c and q . Notice that q cannot be in $\text{int}(D_\rho)$ because $\|p - q\| \geq 2\rho$. Also, there are no other points of S in $Z \subset \text{int}(X)$. Therefore $\text{int}(D_\rho)$ is an open disk of radius ρ centered on B that contains no points of S , contradicting condition 2. \diamond

Lemma 4.4 *Let p, q be two points on B_i , and let X, Y be the two 1-balls on B_i , $\text{bd}(X) = \text{bd}(Y) = \{p, q\}$. If $\text{int}(X) \cap S = \emptyset$ then there exists D_ρ such that $p, q \in \text{bd}(D_\rho)$ and $D_\rho \cap B = X$.*

Proof: Notice that by Lemma 4.2, either $D_\rho \cap B = X$, or $D_\rho \cap B = Y$. It will therefore suffice to show that there must be a point of X other than p, q in D_ρ .

By Lemma 4.3, $\|p - q\| < 2\rho$, and therefore there are two disks $D_{1,\rho}, D_{2,\rho}$ such that $p, q \in \text{bd}(D_{k,\rho}), k = 1, 2$, whose centers lie on the opposite sides of the line through p, q . Assume that there are no points of X other than p, q in either of these disks.

Consider the line through the midpoint of segment p, q and orthogonal to the segment. This line must intersect X at a point c , which lies outside the two disks. It is easy to see that $\|c - p\| = \|c - q\| > \sqrt{2}\rho > \rho$. Then take the disk D_ρ centered in $c \in X$. Since it contains a point of B in its interior, it must intersect B in a 1-ball I containing c , by condition 1. Observe that I cannot include p or q because of the bound on the distance. Therefore, I must be a proper subset of X . Since X does not contain points of S in its interior, $\text{int}(D_\rho)$ violates condition 2. \diamond

Lemma 4.5 *Let p, q be two points on B_i , and let X, Y be the two 1-balls on B_i , $\text{bd}(X) = \text{bd}(Y) =$*

$\{p, q\}$. If $\text{int}(X) \cap S = \emptyset$ then there exist two disks $D_{1,\rho}, D_{2,\rho}$ such that $p, q \in \text{bd}(D_{k,\rho}), k = 1, 2$ and $D_{1,\rho} \cap B = D_{2,\rho} \cap B = X$.

Proof: Let the two disks $D_{k,\rho}$ be as in Lemma 4.4. By that same lemma, one of the two disks, say $D_{1,\rho}$ must be such that $D_{1,\rho} \cap B = X$. Then assume that for the other disk $D_{2,\rho} \cap B \neq X$. By Lemma 4.2 we must have $D_{2,\rho} \cap B = Y$. All of B_i is then contained in the union of the two disks.

Now consider a disk $D_\rho(t)$ centered in $c = tc_1 + (1-t)c_2$, where c_1, c_2 are the centers of $D_{1,\rho}$ and $D_{2,\rho}$, respectively. For $0 \leq t \leq 1$ the disk moves from a position coincident with $D_{1,\rho}$ to one coincident with $D_{2,\rho}$. For each $0 \leq t \leq 1$, $D_\rho(t)$ contains p and q , and therefore, to satisfy condition 1, must contain all X or all Y , but can never contain both.

For any point $x \in \text{int}(X)$ the function

$$f_x(t) = \|x - c(t)\| - \rho$$

is continuous, and negative for $t = 0$. Since $D_\rho(1) \cap \text{int}(X) \neq \text{int}(X)$, there exists $\bar{x} \in \text{int}(X)$ such that $f_{\bar{x}}(1) > 0$. Then there is a $0 < \bar{t} < 1$ such that $f_{\bar{x}}(\bar{t}) = 0$. Let \bar{x} be the point for which is minimum the \bar{t} that makes $f_{\bar{x}}(\bar{t})$ zero.

Then X lies all in $D_\rho(\bar{t})$, and \bar{x} lies on the boundary of $D_\rho(\bar{t})$. Since $\bar{x} \in \text{int}(X)$, and $p, q \in \text{int}(D_\rho(\bar{t}))$, $I = D_\rho(\bar{t}) \cap B$ contains \bar{x} in its interior. But then condition 1 cannot be satisfied. \diamond

Lemma 4.6 *Consider two points $p, q \in S$. If $p \in B_i$ and $q \in B_j, i \neq j$, then the segment $\sigma_T, T = \{p, q\}$ is not α -exposed.*

Proof: For σ_T to be α -exposed there must exist a D_ρ such that $p, q \in \text{bd}(D_\rho)$. But then $D_\rho \cap B$ must be a 1-ball by condition 1, which is impossible since p, q belong to different components of B . \diamond

Lemma 4.7 *Consider two points $p, q \in S$, with $p, q \in B_i$, and let X, Y be the two 1-balls on B_i , $\text{bd}(X) = \text{bd}(Y) = \{p, q\}$. If both $\text{int}(X)$ and $\text{int}(Y)$ contain points of S , then the segment $\sigma_T, T = \{p, q\}$ is not α -exposed.*

Proof: If there exists D_ρ such that $p, q \in \text{bd}(D_\rho)$, then by Lemma 4.2 D_ρ must contain either $\text{int}(X)$ or $\text{int}(Y)$. Since both contain points of S , σ_T cannot be α -exposed. If the disk D_ρ does not exist that σ_T cannot be α -exposed. \diamond

Lemma 4.8 *Consider two points $p, q \in S$, with $p, q \in B_i$, and let X, Y be the two 1-balls on B_i , $\text{bd}(X) = \text{bd}(Y) = \{p, q\}$. If $\text{int}(X) \cap S = \emptyset$, then the segment $\sigma_T, T = \{p, q\}$ is α -exposed. Moreover, σ_T does not bound the interior of \mathcal{S}_α (or, equivalently, σ_T is a singular simplex of the alpha-complex \mathcal{K}_α).*

Proof: By Lemma 4.4 there exist two disks $D_{1,\rho}, D_{2,\rho}$ such that $p, q \in \text{bd}(D_\rho)$ and $D_{k,\rho} \cap B = X, k = 1, 2$. Since $\text{int}(X)$ does not contain points of S , σ_T is α -exposed, and there are two weighted points $x, y, w_x = w_y = \rho^2$ that identify σ_T as α -exposed. \diamond

Lemma 4.9 *There are at least three points of S on each B_i .*

Proof: B_i cannot have 0 points on it, because otherwise condition 2 would be violated for any $\text{int}(D_\rho)$ centered on B_i . Suppose B_i has only one point p of S . Then take D_ρ centered in p . By condition 1, D_ρ intersects B_i in a 1-ball I containing p . Then consider a point $c \in B_i - I$, and a disk D_ρ centered in c . Clearly this disk cannot contain p . Therefore, $\text{int}(D_\rho)$ does not contain any point of S , violating condition 1. For the case of only two points of S on B_i one can repeat the reasoning in Lemma 4.4 and conclude again that condition 1 would not be satisfied. \diamond

We are now ready to prove Theorem 4.1:

Proof: (i) $\mathcal{S}_\alpha(S)$ and B are homeomorphic.

By Lemma 4.9 there are at least three points of S on each connected component B_i of B . For each of these points, say p , there are exactly two other points of S on B_i , say q_1, q_2 , such that the two 1-balls on B_i having p, q_k ($k = 1, 2$) as boundary do not contain any other point of S . Therefore, by Lemmas 4.6-4.7, for each point of S there are exactly two incident 1-simplices in \mathcal{S}_α . Observe that these segments cannot intersect each other in their interior. This could be easily proved here, but it will suffice to notice that the segments are part of the 1-skeleton of a simplicial complex. The α -exposed segments form a one 1-sphere for each component of B . We can then build a homeomorphism by mapping each segment $\sigma_T, T = \{p, q\}$ to the 1-ball $X \subset B_i$ that has p, q as boundary points and contains no other points of S .

(ii) $D(\mathcal{S}_\alpha(S), B) < \rho$.

Each segment $\sigma_T, T = \{p, q\}, p, q \in B_i$ of \mathcal{S}_α is mapped by the homeomorphism to a 1-ball $X \subset B_i$. This ball, by Lemma 4.5, is contained in the intersection of the two disks $D_{1,\rho}, D_{2,\rho}, p, q \in \text{bd}(D_{k,\rho}), k = 1, 2$ (see Figure 5). It is easy to see that for a point x in this intersection, the maximum distance δ to the closest point on the segment σ_T is $\delta < \rho$. Since this is true for all segments of \mathcal{S}_α , the bound holds. \diamond

Notice that locally the error bound can be made arbitrarily small. In fact, for each segment $\sigma_T, T = \{p, q\}$, if $\|p - q\| = 2d$, the maximum local error is

$$\delta < \rho - \sqrt{\rho^2 - d^2}$$




FIGURE 5: The maximum distance δ of a point x on the segment p, q to the closest point of B is bounded by ρ .

which has limit zero as d tends to zero.

Therefore, while a ρ -dense sampling will suffice to reconstruct the manifold B with distance bounded by ρ , we can always make the approximation error arbitrarily small in any region $C \subseteq B$ by simply sampling C at a higher density. Also note that the expression for δ converges to zero quadratically, that is it is sufficient to double the density of the sampling to reduce the error by a factor of four.

We are currently working on extending the theorem above to the 3D case, as well as to the more general case of weighted points in \mathbf{R}^n . We state the 3D version of the theorem here as a

Conjecture 4.1 *Let $B \subset \mathbf{R}^3$ be a compact 1-manifold without boundary, and $S \subset B$ a finite point set. If*

1. *For any closed ball $D_\rho \subset \mathbf{R}^3$ of radius ρ , $B \cap D_\rho$ is either (a) empty; (b) a single point p (then $p \in \text{bd}(D_\rho)$); (c) homeomorphic to a closed 2-ball I , such that $\text{int}(D_\rho) \cap B = \text{int}(I)$;*
2. *An open ball of radius ρ centered on B contains at least one point of S ,*

then the alpha-shape \mathcal{S}_α of S , $\alpha = \rho^2$ is homeomorphic to B and

$$D(\mathcal{S}_\alpha, B) = \max_{p \in \mathcal{S}_\alpha} \min_{q \in B} \|p - q\| < \rho.$$

The conditions above restrict the domain of applicability of our reconstruction tool to surfaces whose




FIGURE 6: A small neighborhood of regions of curvature higher than ρ can be incorrectly reconstructed by the alpha-shape \mathcal{S}_{ρ^2} . Bold segments represent “extraneous” alpha-exposed 1-simplices. (a) A convex sharp feature and a concave high-curvature feature. (b) Extraneous alpha-exposed 1-simplex (detail).

radius of curvature is larger than ρ , as otherwise the ball-intersection requirement would be impossible to satisfy (see Figure 6). Note however the following: (i) This restriction parallels the band-limited requirement in Nyquist’s sampling theorem; (ii) ρ can be made (at least in theory) arbitrarily small. The price to pay to reconstruct small-scale features is to use a high-density sampling, which is reasonable. On a more practical side: (iii) the sampling density of laser scanners is usually much smaller than object features of interest (otherwise large measurement errors would occur); (iv) points are not sampled on the sharp feature, but in its proximity; and (v) data collected in proximity of sharp (or high-curvature) features is usually subject to noise, and therefore not reliable. Accurately reconstructing sharp features (for example to segment the surface into a collection of smooth faces) requires an elaborate analysis of the data and/or additional knowledge of surface characteristics.

5 Conclusions

While the theorems above give us sufficient conditions for a sampling to allow a faithful reconstruction using α -shapes, in practice one has often to deal with less than ideal scans.

In general, i.e. when the conditions of the theorems above are not satisfied, an alpha-shape is a non-connected, mixed-dimension polytope. We are interested in reconstructing *solids*, and therefore it is convenient to define a “regularized” version of an alpha-

shape. The regularization should eliminate dangling and isolated faces, edges, and points from the alpha-shape, and recognize solid components.

In [2], we define a regularized alpha-solid, and describe an automatic method for the selection of an optimal α value, and a heuristic to improve the resulting approximate reconstruction in areas of insufficient sampling density.

The examples shown in the following have been computed with our automatic selection strategy and alpha-solid improvement technique.

Figures 7, 8 and 9 and Table 1 illustrate some examples of alpha-solids computed with the technique described above.

The mechanical part shown in Figure 7(a)-(c) was designed in a commercial solid modeler and randomly sampled.

The object in Figure 7(d)-(f) (courtesy of Jörg Peters, Purdue University) was constructed with a subdivision smoothing of a polyhedron (the result of the union of three polyhedral approximation of tori, with their axis aligned to the coordinate axis). This object has topological genus seven.

The human knee in Figure 7(g)-(i) is a reconstruction from the Visible Human Project data. An isosurface was extracted from the CT volume. As a preprocessing, we reduced the number of vertices from $3 \cdot 10^5$ to about $3 \cdot 10^4$, by replacing clusters of very close points with only one representative (these clusters occur frequently in marching-cube surface extraction).

The data for the golf club and the bunny in Fig-

<i>Object</i>	<i>Number of Points</i>	<i>α-Solid Time</i>	<i>Number of tetrahedra</i>	<i>Removed tetrahedra</i>	<i>Number of Triangles</i>
Femur	9807	1.5	36182	3704	19610
Tibia	9200	1.4	33232	2172	18396
Fibula	8146	1.1	30876	2896	16288
Patella	2050	0.3	7536	683	4096
Part 1	13040	2.5	42507	2473	26088
Club	16864	4.1	58657	754	33142
3 Tori	10833	2.2	42970	2914	21692
Bunny	33123	19.6	127607	3761	66224
Mannequin	10392	2.1	35383	2077	19802

TABLE 1: Results of alpha-solid reconstruction. The table show for each object, from left to right: (1) The number of points in the sampling; (2) The time, in minutes, required by the alpha-solid computation (including 3D Delaunay triangulation, computation of family of alpha-shapes, automatic selection of alpha value, improvement by local sculpturing). All computations were carried out on a SGI Indigo2, with a 250MHz MIPS 4400 CPU; (3) The number of tetrahedra in the initial alpha-solid; (4) The number of tetrahedra removed by the heuristic; (5) The number of triangles in the boundary of the final reconstructed model.

ure 8 was obtained with a laser 3D digitizer.

Figure 9 illustrates the use of weighted alpha-shapes to reconstruct objects that have been sampled at multiple resolution. Some parts of the mannequin head were scanned at a relatively coarse resolution, and other more complicated parts at a finer resolution. Appropriate weights were assigned to the points of each scan.

Table 1 summarizes results and timings on the examples shown.

Once a triangle mesh has been constructed from the data points, one can apply mesh simplification, and subdivision [10], parametric [6, 12] or implicit [1, 2] patch fitting.

We are currently working on a proof for the 3D and general-dimension, weighted points version of the sampling theorem.

Other directions for further research include efficient methods for the computation of two-manifold alpha-shape from the data points without computing the 3D Delaunay (or regular for the weighted case) triangulation. It would also be useful to develop a “real-time”, incremental reconstruction methodology. With this approach, the partially reconstructed surface would be shown to the user as points get scanned.

Acknowledgments. Valerio Pascucci and Guglielmo Rabbio provided useful comments on the sampling theorem.

Thanks to the Computer Graphics Group, Univer-

sity of Washington and the Stanford University Computer Graphics Laboratory, for making some of the range data used in this paper publicly available.

References

- [1] BAJAJ, C., BERNARDINI, F., AND XU, G. Automatic reconstruction of surfaces and scalar fields from 3D scans. In *Computer Graphics Proceedings (1995), Annual Conference Series*. Proceedings of SIGGRAPH 95, ACM SIGGRAPH, pp. 109–118.
- [2] BERNARDINI, F., BAJAJ, C., CHEN, J., AND SCHIKORE, D. Automatic reconstruction of 3D CAD models from digital scans. Submitted for publication, 1996.
- [3] BOISSONNAT, J.-D. Geometric structures for three-dimensional shape representation. *ACM Trans. Graph.* 3, 4 (1984), 266–286.
- [4] CHOI, B. K., SHIN, H. Y., YOON, Y. I., AND LEE, J. W. Triangulation of scattered data in 3D space. *Computer Aided Design* 20, 5 (June 1988), 239–248.
- [5] CURESS, B., AND LEVOY, M. A volumetric method for building complex models from range images. In *Computer Graphics Proceedings (1996), Annual Conference Series*. Proceedings of SIGGRAPH 96, ACM SIGGRAPH, pp. 303–312.
- [6] ECK, M., AND HOPPE, H. Automatic reconstruction of B-splines surfaces of arbitrary topological type. In *Computer Graphics Proceedings (1996), Annual Conference Series*. Proceedings of SIGGRAPH 96, ACM SIGGRAPH, pp. 325–334.
- [7] EDELSBRUNNER, H. Weighted alpha shapes. Tech. Rep. UIUCDCS-R-92-1760, Department of Computer Science, University of Illinois, Urbana-Champaign, IL, 1992.




FIGURE 7: Examples of alpha-solids. (a) The data points are from a random sampling of a model (Part 1) created with a commercial solid modeler. (b) Reconstructed alpha-solid. (c) After improvement by local sculpturing. (d)-(f) The 3-Tori point set is a random sampling of an object created with a solid modeler. (h)-(i) The knee data is from an isosurface extracted from the Visible Human Project data set.

- [8] EDELSBRUNNER, H., KIRKPATRICK, D. G., AND SEIDEL, R. On the shape of a set of points in the plane. *IEEE Trans. Inform. Theory IT-29* (1983), 551–559.
- [9] EDELSBRUNNER, H., AND MÜCKE, E. P. Three-dimensional alpha shapes. *ACM Trans. Graph.* *13*, 1 (Jan. 1994), 43–72.
- [10] HOPPE, H., DEROSE, T., DUCHAMP, T., HALSTEAD, M., JIN, H., McDONALD, J., SCHWITZER, J., AND STUELZLE, W. Piecewise smooth surface reconstruction. In *Computer Graphics Proceedings* (1994), Annual Conference Series. Proceedings of SIGGRAPH 94, ACM SIGGRAPH, pp. 295–302.
- [11] HOPPE, H., DEROSE, T., DUCHAMP, T., McDONALD, J., AND STUELZLE, W. Surface reconstruction from unorganized points. *Computer Graphics* *26*, 2 (July 1992), 71–78. Proceedings of SIGGRAPH 92.
- [12] KRISHNAMURTHY, V., AND LEVOY, M. Fitting smooth surfaces to dense polygonal meshes. In *Computer Graphics Proceedings* (1996), Annual Conference Series. Proceedings of SIGGRAPH 96, ACM SIGGRAPH, pp. 313–324.
- [13] LORENSEN, W., AND CLINE, H. Marching cubes: A high resolution 3D surface construction algorithm. *Computer Graphics* *21* (1987), 163–169.
- [14] TURK, G., AND LEVOY, M. Zippered polygonal meshes from range images. In *Computer Graphics Proceedings* (1994), Annual Conference Series. Proceedings of SIGGRAPH 94, ACM SIGGRAPH, pp. 311–318.
- [15] VELTKAMP, R. C. *Closed object boundaries from scattered points*. PhD thesis, Center for Mathematics and Computer Science, Amsterdam, 1992.

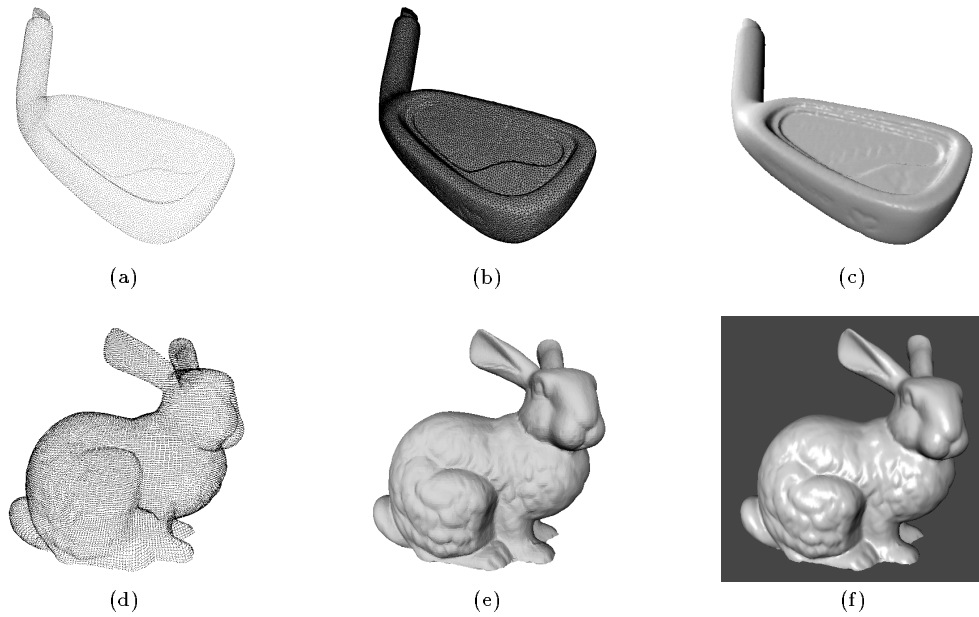


FIGURE 8: Reconstruction from range data. (a) and (d) Combined scans. (b) and (e) Reconstructed alpha-solid. (c) and (f) Phong-shaded rendering.

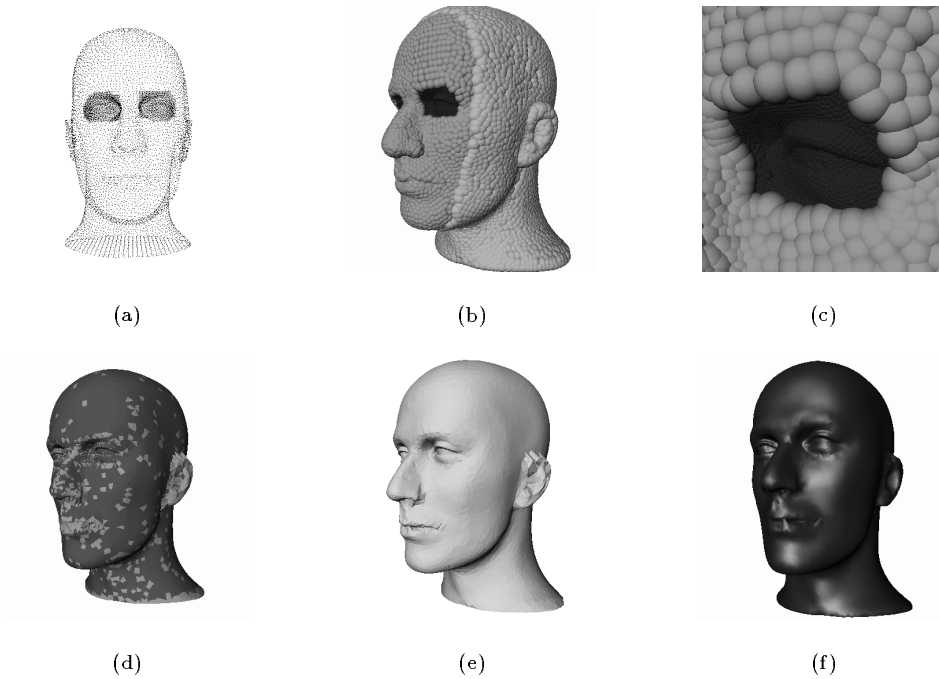


FIGURE 9: Example of reconstruction from a multi-resolution scan using weighted alpha-shapes. (a) Sampling. Notice how the eyes area has been scanned at higher resolution. (b) and (c) Weighted points represented as balls. Weights were assigned manually to simulate a multi-resolution scan. (d) Reconstructed alpha-solid. (e) Alpha-solid after improvement by sculpturing. (f) The same reconstructed model Phong-shaded.